Document Database Service

User Guide

Issue 01

Date 2025-02-13

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

Notice

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

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

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to the *Vul. Response Process*. For details about this process, visit the following web page:
https://www.huawei.com/en/psirt/vul-response-process
For vulnerability information, enterprise customers can visit the following web page:
https://securitybulletin.huawei.com/enterprise/en/security-advisory

Contents

1 Buying an Instance........ 1
1.1 Buying a Cluster Instance........ 1
1.1.1 Quick Config........ 1
1.1.2 Custom Config........ 7
1.2 Buying a Replica Set Instance........ 19
1.2.1 Quick Config........ 19
1.2.2 Custom Config........ 24

2 Connecting to a DB Instance........ 37
2.1 Connecting to a Cluster Instance........ 37
2.1.1 Connection Methods........ 37
2.1.2 (Recommended) Connecting to Cluster Instances Through DAS........ 38
2.1.2.1 Overview........ 38
2.1.2.2 Connecting to a Cluster Instance Through DAS........ 39
2.1.3 Connecting to a Cluster Instance over a Private Network........ 39
2.1.3.1 Configuring Security Group Rules........ 39
2.1.3.2 Connecting to a Cluster Instance Using Mongo Shell (Private Network)........ 43
2.1.3.3 Connecting to Read Replicas Using Mongo Shell........ 53
2.1.4 Connecting to a Cluster Instance over a Public Network........ 57
2.1.4.1 Binding and Unbinding an EIP........ 57
2.1.4.2 Configuring a Security Group........ 60
2.1.4.3 Connecting to a Cluster Instance Using Mongo Shell (Public Network)........ 63
2.1.4.4 Connecting to a Cluster Instance Using Robo 3T........ 70
2.1.5 Connecting to a Cluster Instance Using Program Code........ 77
2.1.5.1 Java........ 77
2.1.5.2 Python........ 81
2.1.5.3 PHP........ 81
2.2 Connecting to a Replica Set Instance........ 83
2.2.1 Connection Methods........ 83
2.2.2 (Recommended) Connecting to Replica Set Instances Through DAS........ 84
2.2.2.1 Overview........ 84
2.2.2.2 Connecting to a Replica Set Instance Through DAS........ 85
2.2.3 Connecting to a Replica Set Instance over a Private Network........ 85
2.2.3.1 Configuring Security Group Rules........ 85

2.2.3.2 Connecting to a Replica Set Instance Using Mongo Shell (Private Network)..89
2.2.3.3 Connecting to Read Replicas Using Mongo Shell.. 101
2.2.4 Connecting to a Replica Set Instance over a Public Network.. 105
2.2.4.1 Binding an EIP...105
2.2.4.2 Configuring Security Group Rules...108
2.2.4.3 Connecting to a Replica Set Instance Using Mongo Shell (Public Network)....................................... 111
2.2.4.4 Connecting to a Replica Set Instance Using Robo 3T.. 118
2.2.5 Connecting to a Replica Set Instance Using Program Code...124
2.2.5.1 Java.. 124
2.2.5.2 Python... 128
2.2.5.3 PHP... 128
2.3 Connecting to a Single Node Instance..130
2.3.1 Connection Methods..130
2.3.2 (Recommended) Connecting to a Single Node Instance Through DAS... 131
2.3.2.1 Overview...131
2.3.2.2 Connecting to a Single Node Instance Through DAS.. 132
2.3.3 Connecting to a Single Node Instance over a Private Network... 132
2.3.3.1 Configuring a Security Group... 133
2.3.3.2 Connecting to a Single Node Instance Using Mongo Shell (Private Network).................................. 136
2.3.4 Connecting to a Single Node Instance over a Public Network.. 140
2.3.4.1 Binding an EIP...140
2.3.4.2 Configuring a Security Group... 143
2.3.4.3 Connecting to a Single Node Instance Using Mongo Shell (Public Network)......................................146
2.3.4.4 Connecting to a Single Node Instance Using Robo 3T...150
2.3.5 Connecting to a Single Node Instance Using Program Code... 157
2.3.5.1 Java.. 157
2.3.5.2 Python... 160
2.3.5.3 PHP... 161
3 Data Migration.. 164
3.1 Migration Scheme Overview... 164
3.2 Migrating Data Using DRS.. 165
3.3 Migrating Data Using mongoexport and mongoimport... 168
3.4 Migrating Data Using mongodump and mongorestore...171
4 Performance Tuning... 176
4.1 Parameters..176
4.2 Read and Write Performance...177
4.3 High CPU Usage..178
4.4 High Storage Usage..184
4.5 High Memory Usage... 186
4.6 Load Imbalance of Cluster Instances..187
4.7 Slow Request Locating... 189
4.8 Statement Optimization.. 191

4.9 Sharding 194

5 Permissions Management 199

5.1 Creating a User and Granting Permissions 199

5.2 Creating a Custom Policy 200

5.3 Syntax of RBAC Policies 202

6 Instance Lifecycle Management 205

6.1 Instance Statuses 205

6.2 Exporting Instance Information 206

6.3 Restarting an Instance or a Node 209

6.4 Deleting a Pay-per-Use Instance 214

6.5 Recycling an Instance 215

6.5.1 Modifying the Recycling Policy 215

6.5.2 Rebuilding an Instance 216

7 Instance Modifications 218

7.1 Changing an Instance Name 218

7.2 Changing an Instance Description 218

7.3 Modifying an Instance Tag 219

7.4 Changing the Name of the Replica Set in the Connection Address 220

7.5 Upgrading a Minor Engine Version 220

7.6 Upgrading a Major Engine Version 223

7.7 Scaling Up Storage Space 224

7.7.1 Scaling Up a Cluster Instance 224

7.7.2 Scaling Up a Replica Set Instance 227

7.7.3 Scaling Up a Read Replica 230

7.7.4 Scaling Up a Single Node Instance 232

7.8 Changing an Instance Class 234

7.8.1 Changing a Cluster Instance Class 234

7.8.2 Changing a Replica Set Instance Class 243

7.8.3 Changing a Single Node Instance Class 248

7.9 Changing Cluster Instance Nodes 250

7.9.1 Adding Cluster Instance Nodes 250

7.9.2 Adding Read Replicas to a Cluster Instance 253

7.9.3 Manually Switching the Primary and Secondary Nodes of a Cluster 255

7.9.4 Reverting Cluster Instance Nodes 259

7.10 Changing Replica Set Instance Nodes 259

7.10.1 Adding Replica Set Instance Nodes 259

7.10.2 Adding Read Replicas to a Replica Set Instance 261

7.10.3 Manually Switching the Primary and Secondary Nodes of a Replica Set 263

7.10.4 Deleting Replica Set Instance Nodes 265

7.10.5 Deleting Read Replicas from a Replica Set Instance 267

7.11 Configuring the Maintenance Window 269

7.12 Changing an AZ........271
7.13 Updating the OS of a DB Instance........273
8 Data Backups........274
8.1 Backup Principles and Solutions........274
8.2 Configuring an Automated Backup Policy........277
8.3 Configuring an Incremental Backup Policy........285
8.4 Configuring the Cross-Region Backup Policy........288
8.5 Setting Backup Method for a DB Instance........292
8.6 Creating a Manual Backup........294
8.7 Deleting a Manual Backup........296
8.8 Deleting an Automated Backup........297
8.9 Stopping a Backup........299
8.10 Downloading a Backup File........300
8.10.1 Using OBS Browser+........300
8.10.2 Using Current Browser........302
8.10.3 Using Download URL........303
9 Data Restorations........306
9.1 Solutions........306
9.2 Restoring Data to a New Instance........307
9.2.1 Restoring a Cluster Backup to a New Instance........307
9.2.2 Restoring a Replica Set Backup to a New Instance........309
9.2.3 Restoring a Single Node Backup to a New Instance........310
9.2.4 Restoring a Cross-Region Backup to a New DB Instance........312
9.3 Restoring Data to the Original Instance........313
9.3.1 Restoring a Cluster Backup to the Original Instance........313
9.3.2 Restoring a Replica Set Backup to the Original Instance........315
9.3.3 Restoring a Single Node Backup to the Original Instance........316
9.4 Restoring Data to a Point in Time........318
9.4.1 Restoring a Cluster Instance to a Point in Time........318
9.4.2 Restoring a Replica Set Instance to a Point in Time........321
9.4.3 Restoring a Replica Set Database and Table to a Point in Time........323
9.4.4 Restoring a Cross-Region Backup to a Point in Time........327
9.5 Restoring Data to an On-Premises Database........329
9.5.1 Restoring a Cluster Backup to an On-premises Database........329
9.5.1.1 Overview........329
9.5.1.2 Directories and Configurations........330
9.5.1.3 Restoring the configsvr Replica Set........332
9.5.1.4 Restoring the shardsvr1 Replica Set........335
9.5.1.5 Restoring the shardsvr2 Replica Set........338
9.5.1.6 Restoring the dds mongos Node........341
9.5.1.7 Checking the Cluster Status........342
9.5.2 Restoring a Replica Set Backup to an On-Premises Database........342

9.5.3 Restoring a Single Node Backup to an On-Premises Database....345
9.6 Restoring Data of Enhanced Edition....346

10 Parameter Template Management.... 347
10.1 Overview.... 347
10.2 Creating a Parameter Template.... 348
10.3 Modifying DDS DB Instance Parameters.... 350
10.4 Viewing Parameter Change History.... 353
10.5 Exporting a Parameter Template....354
10.6 Comparing Parameter Templates....355
10.7 Replicating a Parameter Template....356
10.8 Resetting a Parameter Template.... 357
10.9 Applying a Parameter Template.... 358
10.10 Viewing Application Records of a Parameter Template....359
10.11 Modifying the Description of a Parameter Template.... 359
10.12 Deleting a Parameter Template.... 360

11 Connection Management.... 361
11.1 Querying DB Instance Connections and Managing Sessions....361
11.2 Configuring Cross-CIDR Access.... 365
11.3 Enabling IP Addresses of Shard and Config Nodes.... 366
11.4 Changing a Private IP Address....373
11.5 Changing a Database Port....375
11.6 Applying for and Modifying a Private Domain Name.... 376

12 Database Usage....379
12.1 Creating a Database Account Using Commands.... 379
12.2 Creating a Database Using Commands.... 382
12.3 Which Commands are Supported or Restricted by DDS?.... 383

13 Data Security.... 392
13.1 Enabling or Disabling SSL.... 392
13.2 Resetting the Administrator Password.... 396
13.3 Changing a Security Group....398

14 Monitoring and Alarm Reporting.... 401
14.1 DDS Metrics.... 401
14.2 Configuring Monitoring by Seconds.... 422
14.3 Viewing DDS Metrics.... 424
14.4 Configuring Alarm Rules.... 426
14.5 Managing Alarm Rules....427
14.6 Event Monitoring.... 428
14.6.1 Introduction to Event Monitoring.... 428
14.6.2 Viewing Event Monitoring Data.... 428
14.6.3 Creating an Alarm Rule to Monitor an Event....429

14.6.4 Events Supported by Event Monitoring 432

15 Auditing 437

15.1 Key Operations Recorded by CTS 437

15.2 Viewing CTS Traces 440

16 Logs 441

16.1 Log Reporting 441

16.2 Error Logs 447

16.2.1 Viewing Error Logs on the LTS Console 447

16.2.2 Viewing Error Logs on the DDS Console 449

16.3 Slow Query Logs 451

16.3.1 Viewing Slow Query Logs on the LTS Console 451

16.3.2 Viewing Slow Query Logs on the DDS Console 453

16.4 Audit Logs 458

16.4.1 Audit Log Policy Management 458

16.4.2 Viewing Audit Logs on the LTS Console 463

16.4.3 Viewing Audit Logs on the DDS Console 465

17 Task Center 467

18 DBA Assistant 470

18.1 Real-Time Diagnosis 470

18.1.1 Real-Time Sessions 470

19 SQL Execution Control 472

20 Cross-AZ Disaster Recovery 481

20.1 Creating a Cross-AZ Cluster Instance 481

20.2 Creating a Cross-AZ Replica Set Instance 482

21 Tags 485

21.1 Adding or Modifying a Tag 485

21.2 Filtering Instances by Tag 487

21.3 Deleting a Tag 488

22 Quotas 490

23 DDS Usage Suggestions 492

23.1 Design Rules 492

23.2 Development Rules 494

1 Buying an Instance

1.1 Buying a Cluster Instance

1.1.1 Quick Config

This section describes how to quickly purchase a cluster instance on the management console. DDS helps you quickly configure and create a cluster within several minutes.

Precautions

Each account can create up to 10 cluster instances.

Prerequisites

- You have **registered a Huawei ID and enabled Huawei Cloud services**.
- To display whether the disk is encrypted in the DB instance list, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket**.

Procedure

Step 1 Go to the **Quick Config** page.

Step 2 On the displayed page, select a billing mode and configure information about your DB instance. Then, click **Next**.

Figure 1-1 Basic configurations

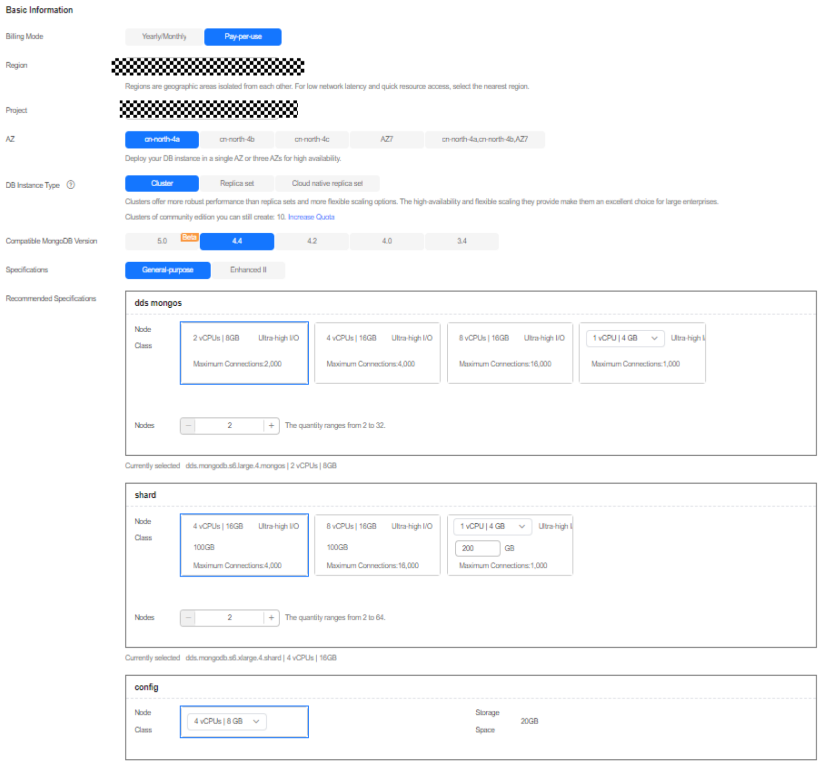

Basic Information

Billing Mode: Yearly/Monthly | Pay-per-use (selected)

Region: [redacted]

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [redacted]

AZ: cn-north-4a (selected) | cn-north-4b | cn-north-4c | AZ7 | cn-north-4a,cn-north-4b,AZ7

Deploy your DB instance in a single AZ or three AZs for high availability.

DB Instance Type: Cluster (selected) | Replica set | Cloud native replica set

Clusters offer more robust performance than replica sets and more flexible scaling options. The high-availability and flexible scaling they provide make them an excellent choice for large enterprises.

Clusters of community edition you can still create: 10. Increase Quota

Compatible MongoDB Version: 5.0 Beta | 4.4 (selected) | 4.2 | 4.0 | 3.4

Specifications: General-purpose (selected) | Enhanced II

Recommended Specifications

dds mongos

Node Class:
- 2 vCPUs | 8GB Ultra-high I/O, Maximum Connections:2,000 (selected)
- 4 vCPUs | 16GB Ultra-high I/O, Maximum Connections:4,000
- 8 vCPUs | 16GB Ultra-high I/O, Maximum Connections:16,000
- 1 vCPU | 4 GB Ultra-high I, Maximum Connections:1,000

Nodes: 2 The quantity ranges from 2 to 32.

Currently selected dds.mongodb.s6.large.4.mongos | 2 vCPUs | 8GB

shard

Node Class:
- 4 vCPUs | 16GB Ultra-high I/O, 100GB, Maximum Connections:4,000 (selected)
- 8 vCPUs | 16GB Ultra-high I/O, 100GB, Maximum Connections:16,000
- 1 vCPU | 4 GB Ultra-high I, 200 GB, Maximum Connections:1,000

Nodes: 2 The quantity ranges from 2 to 64.

Currently selected dds.mongodb.s6.xlarge.4.shard | 4 vCPUs | 16GB

config

Node Class: 4 vCPUs | 8 GB

Storage Space: 20GB

Table 1-1 Basic configurations

Parameter	Description
Billing Mode	Select a billing mode, **Yearly/Monthly** or **Pay-per-use**. ● For yearly/monthly instances – Specify **Required Duration**, and the system deducts the fees incurred from your account based on the service price. – If you do not expect to continue using the instance much after it expires, you can change the billing mode from yearly/monthly to pay-per-use. For details, see **Yearly/ Monthly to Pay-per-Use**. **NOTE** Instances billed on a yearly/monthly basis cannot be deleted. They can only be unsubscribed from. For details, see **Unsubscribing from a Yearly/Monthly Instance**. ● For pay-per-use instances – You are billed for usage based on how much time the service is in use. – If you expect to use the service extensively over a long period of time, you can change its billing mode from pay-per-use to yearly/monthly to reduce costs. For details, see **Pay-per-Use to Yearly/Monthly**.

Parameter	Description
Region	The region where the resource is located. **NOTE** Instances deployed in different regions cannot communicate with each other through a private network, and you cannot change the region of an instance once it is purchased. Exercise caution when selecting a region.
Project	The project corresponds to the current region and can be changed.
AZ	An AZ is a part of a region with its own independent power supply and network. AZs are physically isolated but can communicate through internal network connections. Instances can be deployed in a single AZ or three AZs. **NOTE** The 3-AZ deployment is not available in all regions. If the 3-AZ option is not displayed on the page for you to buy an instance, try a different region. ● If your service requires low network latency between instances, you deploy the components of the instance in the same AZ. If you select a single AZ to deploy your instance, anti-affinity deployment is used by default. With an anti-affinity deployment, your primary, secondary, and hidden nodes are deployed on different physical machines for high availability. ● If you want to deploy an instance across AZs for disaster recovery, select three AZs. In this deployment mode, the dds mongos, shard, and config nodes are evenly distributed across the three AZs.
DB Instance Type	Select **Cluster**. A cluster instance includes three types of nodes: dds mongos, shard, and config. Each shard and config is a three-node replica set to ensure high availability.
Compatible MongoDB Version	● 5.0 ● 4.4 ● 4.2 ● 4.0 ● 3.4

Parameter	Description
CPU Type	DDS supports x86 and Kunpeng CPU architectures. **NOTE** This parameter is available only for MongoDB 4.0 and 3.4. You do not need to set this parameter for other versions. The default value is **x86**. ● x86 x86 CPUs use the Complex Instruction Set Computing (CISC) instruction set. Each instruction can be used to execute low-level hardware operations. CISC instructions vary in length, and tend to be complicated and slow compared to Reduced Instruction Set Computing (RISC). ● Kunpeng The Kunpeng CPU architecture uses RISC. The RISC instruction set is smaller and faster than CISC, thanks to the simplified architecture. Kunpeng CPUs also offer a better balance between power and performance than x86. Kunpeng CPUs offer a high density, low power option that is more cost effective for heavy workloads.
Specifications	With an x86 architecture, you have the following options: ● General-purpose (s6): S6 instances are suitable for applications that require moderate performance generally but occasional bursts of high performance, such as light-workload web servers, enterprise R&D and testing environments, and low- and medium-performance databases. ● Enhanced II (c6): C6 instances have multiple technologies optimized to provide stable powerful compute performance. 25 GE intelligent high-speed NICs are used to provide ultra-high bandwidth and throughput, making it an excellent choice for heavy-load scenarios. It is suitable for websites, web applications, general databases, and cache servers that have higher performance requirements for compute and network resources; and medium- and heavy-load enterprise applications. For details about the supported instance specifications, see **Cluster Instance Specifications**.
dds mongos Node Class	For details about the dds mongos CPU and memory, see **Cluster Instance Specifications**. You can change the class of an instance after it is created. For details, see **Changing the Instance Class**.
dds mongos Nodes	The value ranges from 2 to 32. You can add nodes to an instance after it is created if necessary. For details, see **Adding Cluster Instance Nodes**.

Parameter	Description
shard Node Class	For details about the shard CPU and memory, see **Cluster Instance Specifications**. The shard node stores user data but cannot be accessed directly. You can change the class of an instance after it is created. For details, see **Changing the Instance Class**.
shard Storage Space	The value ranges from 10 GB to 5,000 GB and must be a multiple of 10. You can scale up an instance after it is created. For details, see **Scaling Up a Cluster Instance**. **NOTE** ● If the storage space you purchased exceeds 600 GB and the remaining storage space is 18 GB, the instance becomes **Read-only**. ● If the storage space you purchased is less than 600 GB and the storage space usage reaches 97%, the instance becomes **Read-only**. In these cases, delete unnecessary resources or expand the capacity.
shard Nodes	The value ranges from 2 to 32. You can add nodes to an instance after it is created if necessary. For details, see **Adding Cluster Instance Nodes**.
config Node Class	For details about the CPU and memory of the config node, see **Cluster Instance Specifications**. You can change the class of an instance after it is created. For details, see **Changing the Instance Class**.
config Storage Space	Based on the functions and minimum requirements of the config node, the storage space of the config node is set to 20 GB by default. You cannot scale up the storage of the node after it is created.

Figure 1-2 Network, Required Duration, and Quantity

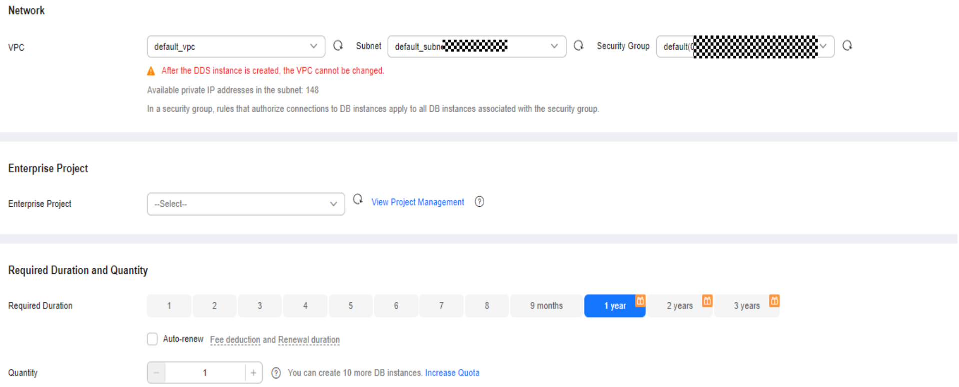

Network

VPC: default_vpc Subnet: default_subne[redacted] Security Group: default([redacted]

After the DDS instance is created, the VPC cannot be changed.

Available private IP addresses in the subnet: 148

In a security group, rules that authorize connections to DB instances apply to all DB instances associated with the security group.

Enterprise Project

Enterprise Project: --Select-- View Project Management

Required Duration and Quantity

Required Duration: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 months | [✓ 1 year] | 2 years | 3 years

□ Auto-renew Fee deduction and Renewal duration

Quantity: 1 You can create 10 more DB instances. Increase Quota

Table 1-2 Network settings

Parameter	Description
VPC	The VPC where your DB instances are located. A VPC isolates networks for different services. It allows you to easily manage and configure private networks and change network configurations. You need to create or select the required VPC. For details, see **Creating a VPC** in the *Virtual Private Cloud User Guide*. For details about the constraints on the use of VPCs, see **Connection Methods**. If there are no VPCs available, DDS creates one for you by default. **NOTE** After the DDS instance is created, the VPC cannot be changed.
Enterprise Project	Only enterprise users can use this function. To use this function, contact customer service. An enterprise project is a cloud resource management mode, in which cloud resources and members are centrally managed by project. Select an enterprise project from the drop-down list. The default project is **default**. For more information about enterprise project, see **Project Management** in *Enterprise Management User Guide*. To customize an enterprise project, click **Enterprise** in the upper right corner of the console. The **Enterprise Management** page is displayed. For details, see **Creating an Enterprise Project** in *Enterprise Management User Guide*.

Table 1-3 Required duration and quantity

Parameter	Description
Required Duration	The length of your subscription if you select **Yearly/Monthly** billing. Subscription lengths range from one month to three years.
Auto-renew	● By default, this option is not selected. ● If you select this option, the auto-renew cycle is determined by the length of the subscription.
Quantity	The purchase quantity depends on the cluster instance quota. If your current quota does not allow you to purchase the required number of instances, you can apply for an increased quota. Yearly/Monthly instances that were purchased in batches have the same specifications except for the instance name and ID.

Step 3 On the displayed page, confirm the instance details.

- For yearly/monthly instances

   - If you need to modify the specifications, click **Previous** to return to the previous page.
   - If you do not need to modify the specifications, read and agree to the service agreement and click **Pay Now** to go to the payment page and complete the payment.
- For pay-per-use instances
   - If you need to modify the specifications, click **Previous** to return to the previous page.
   - If you do not need to modify the specifications, read and agree to the service agreement and click **Submit** to start creating the instance.

Step 4 Click **Back to Instance List**. After a DDS instance is created, you can view and manage it on the **Instances** page.

- When an instance is being created, the status displayed in the **Status** column is **Creating**. This process takes about 15 minutes. After the creation is complete, the status changes to **Available**.
- DDS enables an automated backup policy by default, but you can disable it after an instance is created. An automated full backup is immediately triggered after the creation of an instance.

----End

1.1.2 Custom Config

This section describes how to purchase a cluster instance in custom mode on the management console. You can customize the computing resources and storage space of a cluster instance based on your service requirements. In addition, you can configure advanced settings, such as slow query log and automated backup.

Precautions

Each account can create up to 10 cluster instances.

Prerequisites

- You have **registered a Huawei ID and enabled Huawei Cloud services**.
- To display whether the disk is encrypted in the DB instance list, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket**.

Procedure

Step 1 Go to the **Custom Config** page.

Step 2 On the displayed page, select a billing mode and configure information about your DB instance. Then, click **Next**.

Figure 1-3 Basic configurations

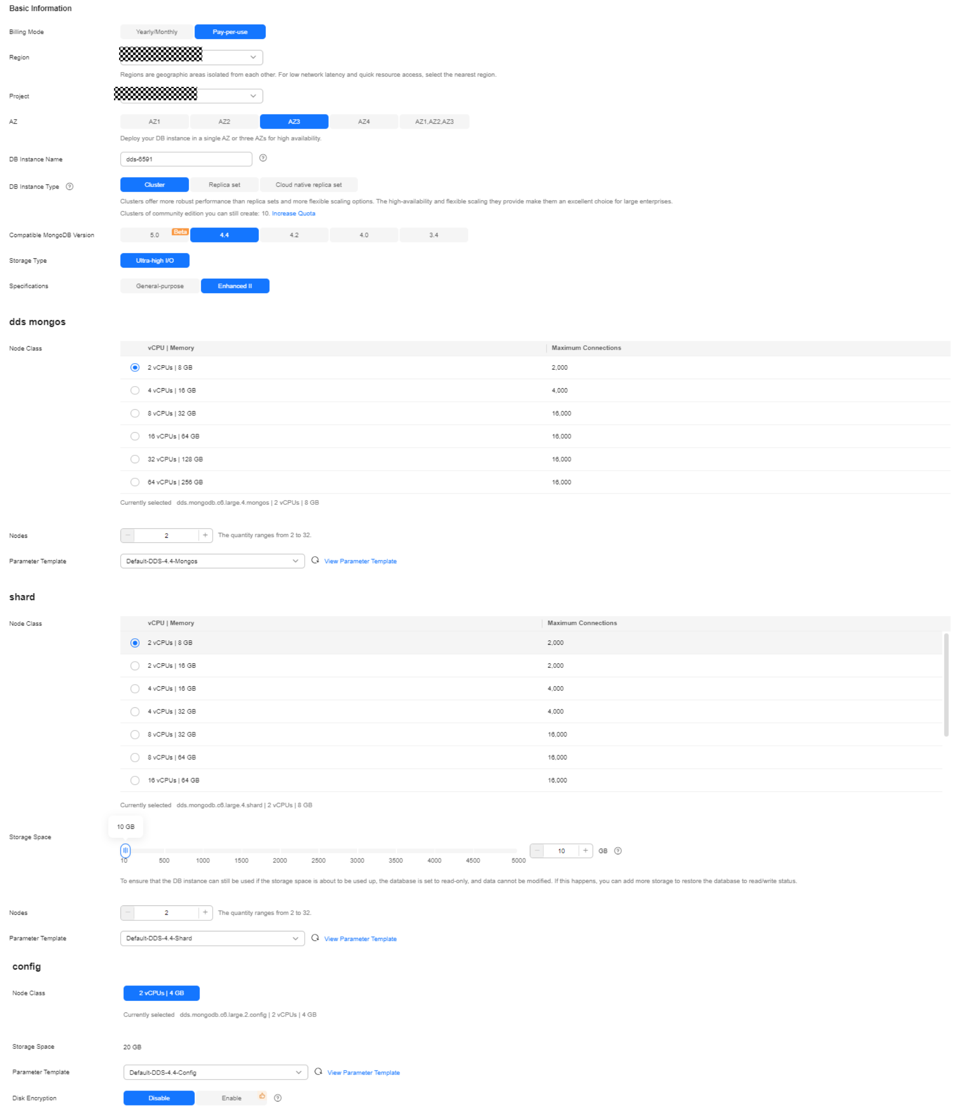

Basic Information

Billing Mode: Yearly/Monthly [✓ Pay-per-use]

Region: [redacted]

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [redacted]

AZ: AZ1 AZ2 [✓ AZ3] AZ4 AZ1,AZ2,AZ3

Deploy your DB instance in a single AZ or three AZs for high availability.

DB Instance Name: dds-6591

DB Instance Type: [✓ Cluster] Replica set Cloud native replica set

Clusters offer more robust performance than replica sets and more flexible scaling options. The high-availability and flexible scaling they provide make them an excellent choice for large enterprises.

Clusters of community edition you can still create: 10. Increase Quota

Compatible MongoDB Version: 5.0 Beta [✓ 4.4] 4.2 4.0 3.4

Storage Type: [✓ Ultra-high I/O]

Specifications: General-purpose [✓ Enhanced II]

dds mongos

Node Class

	vCPU | Memory	Maximum Connections
●	2 vCPUs | 8 GB	2,000
○	4 vCPUs | 16 GB	4,000
○	8 vCPUs | 32 GB	16,000
○	16 vCPUs | 64 GB	16,000
○	32 vCPUs | 128 GB	16,000
○	64 vCPUs | 256 GB	16,000

Currently selected dds.mongodb.c6.large.4.mongos | 2 vCPUs | 8 GB

Nodes: 2 The quantity ranges from 2 to 32.

Parameter Template: Default-DDS-4.4-Mongos View Parameter Template

shard

Node Class

	vCPU | Memory	Maximum Connections
●	2 vCPUs | 8 GB	2,000
○	2 vCPUs | 16 GB	2,000
○	4 vCPUs | 16 GB	4,000
○	4 vCPUs | 32 GB	4,000
○	8 vCPUs | 32 GB	16,000
○	8 vCPUs | 64 GB	16,000
○	16 vCPUs | 64 GB	16,000

Currently selected dds.mongodb.c6.large.4.shard | 2 vCPUs | 8 GB

Storage Space: 10 GB 10 GB

To ensure that the DB instance can still be used if the storage space is about to be used up, the database is set to read-only, and data cannot be modified. If this happens, you can add more storage to restore the database to read/write status.

Nodes: 2 The quantity ranges from 2 to 32.

Parameter Template: Default-DDS-4.4-Shard View Parameter Template

config

Node Class: [✓ 2 vCPUs | 4 GB]

Currently selected dds.mongodb.c6.large.2.config | 2 vCPUs | 4 GB

Storage Space: 20 GB

Parameter Template: Default-DDS-4.4-Config View Parameter Template

Disk Encryption: [✓ Disable] Enable

Table 1-4 Basic configurations

Parameter	Description
Billing Mode	Select a billing mode, **Yearly/Monthly** or **Pay-per-use**. ● For yearly/monthly instances – Specify **Required Duration**, and the system deducts the fees incurred from your account based on the service price. – If you do not expect to continue using the instance much after it expires, you can change the billing mode from yearly/monthly to pay-per-use. For details, see **Yearly/ Monthly to Pay-per-Use**. **NOTE** Instances billed on a yearly/monthly basis cannot be deleted. They can only be unsubscribed from. For details, see **Unsubscribing from a Yearly/Monthly Instance**. ● For pay-per-use instances – You are billed for usage based on how much time the service is in use. – If you expect to use the service extensively over a long period of time, you can change its billing mode from pay-per-use to yearly/monthly to reduce costs. For details, see **Pay-per-Use to Yearly/Monthly**.
Region	The region where the resource is located. **NOTE** Instances deployed in different regions cannot communicate with each other through a private network, and you cannot change the region of an instance once it is purchased. Exercise caution when selecting a region.
Project	The project corresponds to the current region and can be changed.

Parameter	Description
AZ	An AZ is a part of a region with its own independent power supply and network. AZs are physically isolated but can communicate through internal network connections. Instances can be deployed in a single AZ or three AZs. ● If your service requires low network latency between instances, you deploy the components of the instance in the same AZ. If you select a single AZ to deploy your instance, anti-affinity deployment is used by default. With an anti-affinity deployment, your primary, secondary, and hidden nodes are deployed on different physical machines for high availability. ● If you want to deploy an instance across AZs for disaster recovery, select three AZs. In this deployment mode, the dds mongos, shard, and config nodes are evenly distributed across the three AZs. **NOTE** The 3-AZ deployment is not available in all regions. If the 3-AZ option is not displayed on the page for you to buy an instance, try a different region.
DB Instance Name	● The instance name that you specify after the purchase. The instance name must contain 4 to 64 characters and must start with a letter. It is case sensitive and can contain letters, digits, hyphens (-), and underscores (_). It cannot contain other special characters. ● The instance name can be the same as an existing instance name. ● If you buy a batch of instances at once, a 4-digit numerical suffix will be added to the instance names, starting with **-0001**. If you later make another batch purchase, the new instance names will be numbered first using any suffixes missing from the sequence of your existing instances, and then continuing on from where your last batch purchase left off. For example, a batch of 3 instances gets the suffixes **-0001**, **-0002**, and **-0003**. If you deleted instance **0002** and then bought 3 more instances, the new instances would get the suffixes **-0002**, **-0004**, and **-0005**. ● After the DB instance is created, you can change its name. For details, see **Changing an Instance Name**.
DB Instance Type	Select **Cluster**. A cluster instance includes three types of nodes: dds mongos, shard, and config. Each shard and config is a three-node replica set to ensure high availability.

Parameter	Description
Compatible MongoDB Version	● 5.0 ● 4.4 ● 4.2 ● 4.0 ● 3.4
CPU Type	DDS supports x86 and Kunpeng CPU architectures. **NOTE** This parameter is available only for MongoDB 4.0 and 3.4. You do not need to set this parameter for other versions. The default value is **x86**. ● x86 x86 CPUs use the Complex Instruction Set Computing (CISC) instruction set. Each instruction can be used to execute low-level hardware operations. CISC instructions vary in length, and tend to be complicated and slow compared to Reduced Instruction Set Computing (RISC). ● Kunpeng The Kunpeng CPU architecture uses RISC. The RISC instruction set is smaller and faster than CISC, thanks to the simplified architecture. Kunpeng CPUs also offer a better balance between power and performance than x86. Kunpeng CPUs offer a high density, low power option that is more cost effective for heavy workloads.
Storage Type	The storage type can be **Ultra-high I/O** and **Extreme SSD** for non-DeC users. For DeC users, the supported storage types depend on the selected resource type. ● If you select **EVS** for **Resource Type**, **Storage Type** is set to **Cloud SSD**. ● If you select **DSS** for **Resource Type**, **Storage Type** can be set to **Common I/O**, **High I/O**, or **Cloud SSD**.
Storage Engine	● WiredTiger WiredTiger is the default storage engine of DDS 3.4 and 4.0. WiredTiger provides different granularity concurrency control and compression mechanism for data management. It can provide the best performance and storage efficiency for different kinds of applications. ● RocksDB RocksDB is the default storage engine of DDS 4.2 and 4.4. RocksDB supports efficient point lookup, range scan, and high-speed write. RocksDB can be used as the underlying data storage engine of MongoDB and is suitable for scenarios with a large number of write operations.

Parameter	Description
Specifications	With an x86 architecture, you have the following options: ● General-purpose (s6): S6 instances are suitable for applications that require moderate performance generally but occasional bursts of high performance, such as light-workload web servers, enterprise R&D and testing environments, and low- and medium-performance databases. ● Enhanced II (c6): C6 instances have multiple technologies optimized to provide stable powerful compute performance. 25 GE intelligent high-speed NICs are used to provide ultra-high bandwidth and throughput, making it an excellent choice for heavy-load scenarios. It is suitable for websites, web applications, general databases, and cache servers that have higher performance requirements for compute and network resources; and medium- and heavy-load enterprise applications. For details about the supported instance specifications, see **Cluster Instance Specifications**.
dds mongos Node Class	For details about the dds mongos CPU and memory, see **Cluster Instance Specifications**. You can change the class of an instance after it is created. For details, see **Changing the Instance Class**.
dds mongos Nodes	The value ranges from 2 to 32. You can add nodes to an instance after it is created if necessary. For details, see **Adding Cluster Instance Nodes**.
dds mongos Parameter Template	The parameters that apply to the dds mongos nodes. After an instance is created, you can change the parameter template of a node to bring out the best performance. For details, see **Editing a Parameter Template**.
shard Node Class	For details about the shard CPU and memory, see **Cluster Instance Specifications**. The shard node stores user data but cannot be accessed directly. You can change the class of an instance after it is created. For details, see **Changing the Instance Class**.
shard Storage Space	The value ranges from 10 GB to 5,000 GB and must be a multiple of 10. You can scale up an instance after it is created. For details, see **Scaling Up a Cluster Instance**. **NOTE** ● If the storage space you purchased exceeds 600 GB and the remaining storage space is 18 GB, the instance becomes **Read-only**. ● If the storage space you purchased is less than 600 GB and the storage space usage reaches 97%, the instance becomes **Read-only**. In these cases, delete unnecessary resources or expand the capacity.

Parameter	Description
shard Nodes	The value ranges from 2 to 32. You can add nodes to an instance after it is created if necessary. For details, see **Adding Cluster Instance Nodes**.
shard Parameter Template	The parameters that apply to the shard nodes. After an instance is created, you can change the parameter template of a node to bring out the best performance. For details, see **Editing a Parameter Template**.
config Node Class	For details about the CPU and memory of the config node, see **Cluster Instance Specifications**. You can change the class of an instance after it is created. For details, see **Changing the Instance Class**.
config Storage Space	Based on the functions and minimum requirements of the config node, the storage space of the config node is set to 20 GB by default. You cannot scale up the storage of the node after it is created.
config Parameter Template	The parameters that apply to the config nodes. After an instance is created, you can change the parameter template of a node to bring out the best performance. For details, see **Editing a Parameter Template**.
Disk Encryption	● **Disabled**: Disable encryption. ● **Enabled**: Enable encryption. This feature improves data security but slightly affects read/write performance. **Key Name**: Select or create a private key, which is the tenant key. **NOTE** – After an instance is created, the disk encryption status and the key cannot be changed. Disk encryption will not encrypt backup data stored in OBS. To enable backup data encryption, contact customer service. – To check whether the disk is encrypted, you can view **Disk Encrypted** in the DB instance list. – If disk encryption or backup data encryption is enabled, keep the key properly. Once the key is disabled, deleted, or frozen, the database will be unavailable and data may not be restored. If disk encryption is enabled but backup data encryption is not enabled, you can **restore data to a new instance from backups**. If both disk encryption and backup data encryption are enabled, data cannot be restored. – For details about how to create a key, see "**Creating a CMK**" in *Data Encryption Workshop User Guide*. – Disk encryption supports only the AES_256 key algorithm.

Figure 1-4 Administrator settings

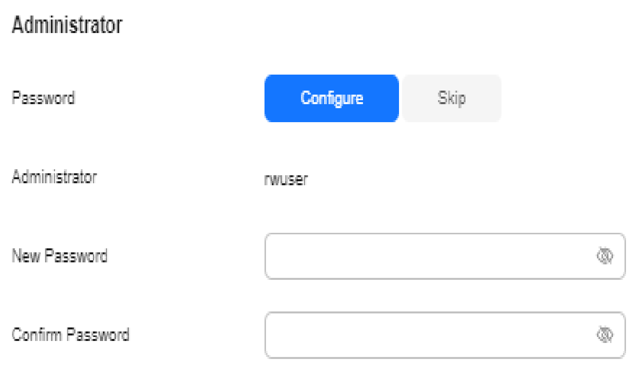

Administrator

Password: Configure | Skip

Administrator: rwuser

New Password

Confirm Password

Table 1-5 Administrator settings

Parameter	Description
Password	● Configure Enter and confirm the new administrator password. After an instance is created, you can connect to the instance using the password. ● Skip To log in, you will have to reset the password later on the **Basic Information** page. If you need to connect to an instance after it is created, locate the instance and choose **More** > **Reset Password** in the **Operation** column to set a password for the instance first.
Administrator	The default account is **rwuser**.
Administrator Password	Set a password for the administrator. The password must be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and at least one of the following special characters: ~!@#%^*-_=+?()$ Keep this password secure. If lost, the system cannot retrieve it for you.
Confirm Password	Enter the administrator password again.

Figure 1-5 Network and required duration

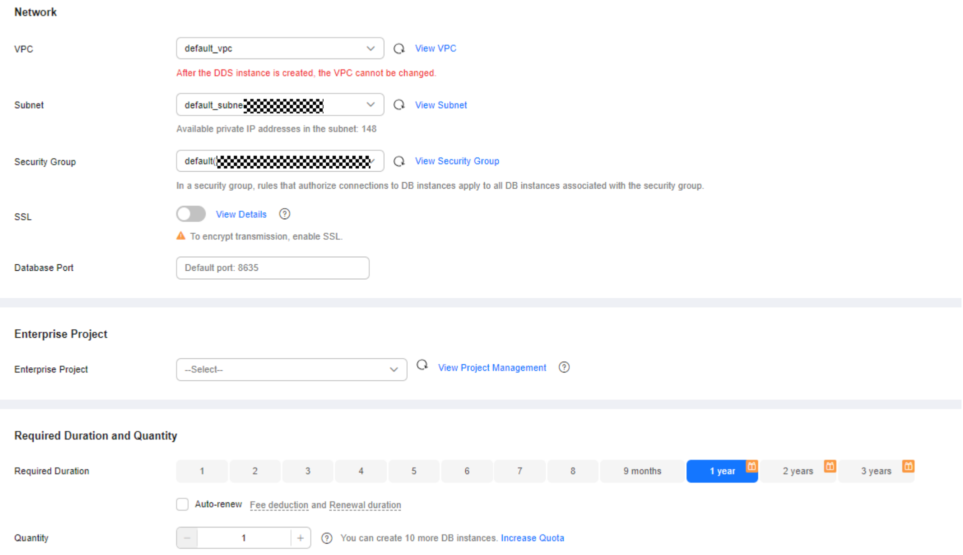

Network

VPC: default_vpc — View VPC

After the DDS instance is created, the VPC cannot be changed.

Subnet: default_subne[redacted] — View Subnet

Available private IP addresses in the subnet: 148

Security Group: default[redacted] — View Security Group

In a security group, rules that authorize connections to DB instances apply to all DB instances associated with the security group.

SSL: View Details

To encrypt transmission, enable SSL.

Database Port: Default port: 8635

Enterprise Project

Enterprise Project: --Select-- — View Project Management

Required Duration and Quantity

Required Duration: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 months | 1 year | 2 years | 3 years

□ Auto-renew Fee deduction and Renewal duration

Quantity: 1 — You can create 10 more DB instances. Increase Quota

Table 1-6 Network settings

Parameter	Description
VPC	The VPC where your DB instances are located. A VPC isolates networks for different services. It allows you to easily manage and configure private networks and change network configurations. You need to create or select the required VPC. For details about how to create a VPC, see "Creating a VPC" in *Virtual Private Cloud User Guide*. For details about the constraints on the use of VPCs, see **Connection Methods**. If there are no VPCs available, DDS creates one for you by default. VPC owners can share the subnets in a VPC with one or multiple accounts through Resource Access Manager (RAM). This allows for more efficient use of network resources and reduces O&M costs. For more information about VPC subnet sharing, see **VPC Sharing** in the *Virtual Private Cloud User Guide*. **NOTE** After the DDS instance is created, the VPC cannot be changed.
Subnet	A subnet provides dedicated network resources that are logically isolated from other networks for security reasons. After the instance is created, you can change the private IP address assigned by the subnet. For details, see **Changing a Private IP Address**. **NOTE** Both IPv4 and IPv6 subnets are supported.

Parameter	Description
Security Group	A security group controls access between DDS and other services. If there are no security groups available, DDS creates one for you by default. **NOTE** ● Ensure that there is a security group rule configured that allows clients to access instances. For example, select an inbound TCP rule with the default port 8635, and enter a subnet IP address or select a security group that the instance belongs to. ● When creating a DB instance, you can select multiple security groups. For better network performance, you are advised to select no more than five security groups. In such a case, the access rules of all the selected security groups apply on the instance.
SSL	Secure Sockets Layer (SSL) encrypts connections between clients and servers, preventing data from being tampered with or stolen during transmission. You can enable SSL to improve data security. After an instance is created, you can connect to it using SSL.
Database Port	The default DDS port is 8635, but this port can be modified if necessary. If you change the port, add a corresponding security group rule to allow access to the instance. **NOTE** ● The database port is the port of the dds mongos node. The default port is 8635. To change the port, see **Changing a Database Port**. ● The shard node port is 8637, and the config node port is 8636, which cannot be changed. For details about how to connect to the shard and config nodes, see **Enabling IP Addresses of Shard and Config Nodes**.
Enterprise Project	Only enterprise users can use this function. To use this function, contact customer service. An enterprise project is a cloud resource management mode, in which cloud resources and members are centrally managed by project. Select an enterprise project from the drop-down list. The default project is **default**. For more information about enterprise project, see *Enterprise Management User Guide*.

Figure 1-6 Advanced settings

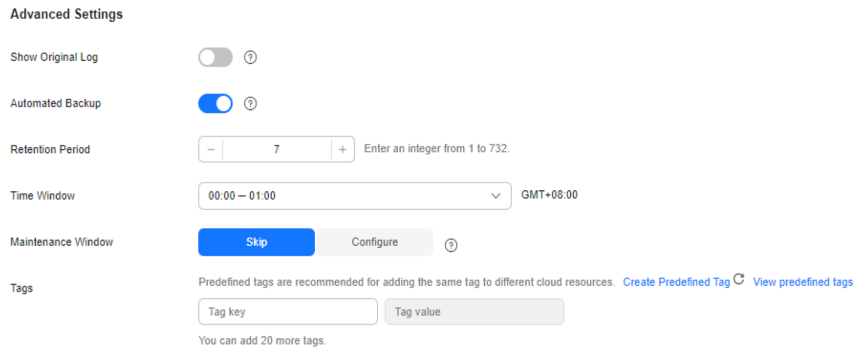

Advanced Settings

Show Original Log: ○ ⓘ

Automated Backup: ● ⓘ

Retention Period: − 7 + Enter an integer from 1 to 732.

Time Window: 00:00 – 01:00 GMT+08:00

Maintenance Window: Skip | Configure ⓘ

Tags: Predefined tags are recommended for adding the same tag to different cloud resources. Create Predefined Tag View predefined tags

Tag key | Tag value

You can add 20 more tags.

Table 1-7 Advanced settings

Parameter	Description
Show Original Log	If Show Original Log is enabled, the original slow query logs will be displayed, and the logs will be transferred to an OBS bucket. By default, the system automatically deletes logs from the OBS bucket after 30 days, and the retention period cannot be changed.
Automated Backup	DDS enables an automated backup policy by default, but you can disable it after an instance is created. An automated full backup is immediately triggered after the creation of an instance. For details, see **Configuring an Automated Backup Policy**.
Retention Period (days)	**Retention Period** refers to the number of days that data is kept. You can increase the retention period to improve data reliability. The backup retention period is from 1 to 732 days.
Time Window	A one-hour period the backup will be scheduled within 24 hours, such as 01:00-02:00. The backup time is in UTC format.

Parameter	Description
Tags	(Optional) You can add tags to DDS instances so that you can quickly search for and filter specified instances by tag. Each DDS instance can have up to 20 tags. If your organization has configured tag policies for DDS, add tags to DB instances based on the policies. If a tag does not comply with the policies, DB instance creation may fail. Contact your organization administrator to learn more about tag policies. ● Create a tag. You can create tags on the DDS console. A tag key and a value are required when you create a tag. Key: This parameter is mandatory. – Each tag key must be unique for each instance. – A tag key consists of up to 36 characters. – The key must consist of only digits, letters, underscores (_), and hyphens (-). Value: This parameter is optional. – The value consists of up to 43 characters. – The value must consist of only digits, letters, underscores (_), periods (.), and hyphens (-). ● Add a predefined tag. Predefined tags can be used to identify multiple cloud resources. To tag a cloud resource, you can select a created predefined tag from the drop-down list, without entering a key and value for the tag. For example, if a predefined tag has been created, its key is Usage and value is Project1. When you configure the key and value for a cloud resource, the created predefined tag will be displayed on the page. After an instance is created, you can click the instance name to view its tags. On the **Tags** page, you can also **modify or delete the tags**. In addition, you can quickly **search for and filter specified instances by tag**. You can add a tag to an instance after the instance is created. For details, see **Adding a Tag**.

If you have any question about the price, click **Price Details**.

NOTE

Instance performance depends on the specifications you select during creation. The hardware configuration items that can be selected include the node class and storage space.

Step 3 On the displayed page, confirm the instance details.

- For yearly/monthly instances
  - If you need to modify the specifications, click **Previous** to return to the previous page.
  - If you do not need to modify the specifications, read and agree to the service agreement and click **Pay Now** to go to the payment page and complete the payment.
- For pay-per-use instances
  - If you need to modify the specifications, click **Previous** to return to the previous page.
  - If you do not need to modify the specifications, read and agree to the service agreement and click **Submit** to start creating the instance.

Step 4 Click **Back to Instance List**. After a DDS instance is created, you can view and manage it on the **Instances** page.

- When an instance is being created, the status displayed in the **Status** column is **Creating**. This process takes about 15 minutes. After the creation is complete, the status changes to **Available**.
- Yearly/Monthly instances that were purchased in batches have the same specifications except for the instance name and ID.

----End

1.2 Buying a Replica Set Instance

1.2.1 Quick Config

This section describes how to quickly purchase a replica set instance on the management console. DDS provides several recommended configurations to help you purchase a replica set instance within several minutes.

Prerequisites

- You have **registered a Huawei ID and enabled Huawei Cloud services**.

Procedure

Step 1 Go to the **Quick Config** page.

Step 2 On the displayed page, select a billing mode and configure information about your DB instance. Then, click **Next**.

Figure 1-7 Basic configurations

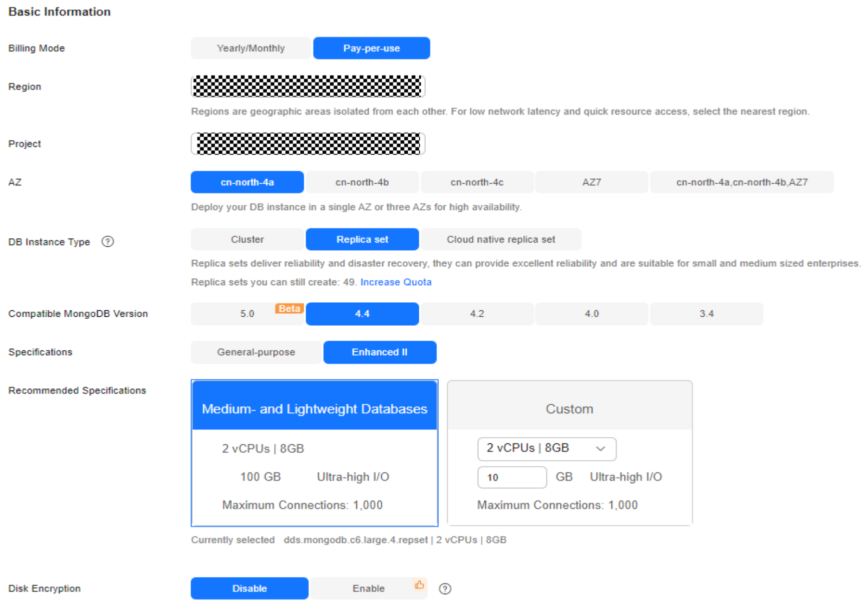

Basic Information

Billing Mode: Yearly/Monthly | **Pay-per-use**

Region: [redacted]

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [redacted]

AZ: **cn-north-4a** | cn-north-4b | cn-north-4c | AZ7 | cn-north-4a,cn-north-4b,AZ7

Deploy your DB instance in a single AZ or three AZs for high availability.

DB Instance Type: Cluster | **Replica set** | Cloud native replica set

Replica sets deliver reliability and disaster recovery, they can provide excellent reliability and are suitable for small and medium sized enterprises.

Replica sets you can still create: 49. Increase Quota

Compatible MongoDB Version: 5.0 Beta | **4.4** | 4.2 | 4.0 | 3.4

Specifications: General-purpose | **Enhanced II**

Recommended Specifications:

Medium- and Lightweight Databases
2 vCPUs | 8GB
100 GB Ultra-high I/O
Maximum Connections: 1,000

Custom
2 vCPUs | 8GB
10 GB Ultra-high I/O
Maximum Connections: 1,000

Currently selected dds.mongodb.c6.large.4.repset | 2 vCPUs | 8GB

Disk Encryption: **Disable** | Enable

Table 1-8 Basic configurations

Parameter	Description
Billing Mode	Select a billing mode, **Yearly/Monthly** or **Pay-per-use**. ● For yearly/monthly instances – Specify **Required Duration**, and the system deducts the fees incurred from your account based on the service price. – If you do not expect to continue using the instance much after it expires, you can change the billing mode from yearly/monthly to pay-per-use. For details, see **Yearly/Monthly to Pay-per-Use**. NOTE Instances billed on a yearly/monthly basis cannot be deleted. They can only be unsubscribed from. For details, see **Unsubscribing from a Yearly/Monthly Instance**. ● For pay-per-use instances – You are billed for usage based on how much time the service is in use. – If you expect to use the service extensively over a long period of time, you can change its billing mode from pay-per-use to yearly/monthly to reduce costs. For details, see **Pay-per-Use to Yearly/Monthly**.
Region	The region where the resource is located. **NOTE** Instances deployed in different regions cannot communicate with each other through a private network, and you cannot change the region of an instance once it is purchased. Exercise caution when selecting a region.

Parameter	Description
Project	The project corresponds to the current region and can be changed.
AZ	An AZ is a part of a region with its own independent power supply and network. AZs are physically isolated but can communicate through internal network connections. Instances can be deployed in a single AZ or three AZs. ● If your service requires low network latency between instances, you deploy the components of the instance in the same AZ. If you select a single AZ to deploy your instance, anti-affinity deployment is used by default. With an anti-affinity deployment, your primary, secondary, and hidden nodes are deployed on different physical machines for high availability. ● If you want to deploy an instance across AZs for disaster recovery, select three AZs. In this deployment mode, the primary, secondary, and hidden nodes are evenly distributed across three AZs. **NOTE** The 3-AZ deployment is not available in all regions. If the 3-AZ option is not displayed on the page for you to buy an instance, try a different region.
DB Instance Type	Select **Replica set**. A replica set consists of the primary node, secondary node, and hidden node. If a primary node goes down or becomes faulty, a secondary node is automatically assigned to the primary role and continues normal operation. If a secondary node is unavailable, a hidden node will take the role of the secondary to ensure high availability.
Compatible MongoDB Version	● 5.0 ● 4.4 ● 4.2 ● 4.0 ● 3.4

Parameter	Description
CPU Type	DDS supports x86 and Kunpeng CPU architectures. **NOTE** This parameter is available only for MongoDB 4.0 and 3.4. You do not need to set this parameter for other versions. The default value is **x86**. ● x86 x86 CPUs use the Complex Instruction Set Computing (CISC) instruction set. Each instruction can be used to execute low-level hardware operations. CISC instructions vary in length, and tend to be complicated and slow compared to Reduced Instruction Set Computing (RISC). ● Kunpeng The Kunpeng CPU architecture uses RISC. The RISC instruction set is smaller and faster than CISC, thanks to the simplified architecture. Kunpeng CPUs also offer a better balance between power and performance than x86. Kunpeng CPUs offer a high density, low power option that is more cost effective for heavy workloads.
Specifications	With an x86 architecture, you have the following options: ● General-purpose (s6): S6 instances are suitable for applications that require moderate performance generally but occasional bursts of high performance, such as light-workload web servers, enterprise R&D and testing environments, and low- and medium-performance databases. ● Enhanced II (c6): C6 instances have multiple technologies optimized to provide stable powerful compute performance. 25 GE intelligent high-speed NICs are used to provide ultra-high bandwidth and throughput, making it an excellent choice for heavy-load scenarios. It is suitable for websites, web applications, general databases, and cache servers that have higher performance requirements for compute and network resources; and medium- and heavy-load enterprise applications.
Recommended Specifications	Currently, medium- and lightweight database specifications and customized specifications are supported. **NOTE** ● The storage space ranges from 10 GB to 5,000 GB.

Figure 1-8 Network, Required Duration, and Quantity

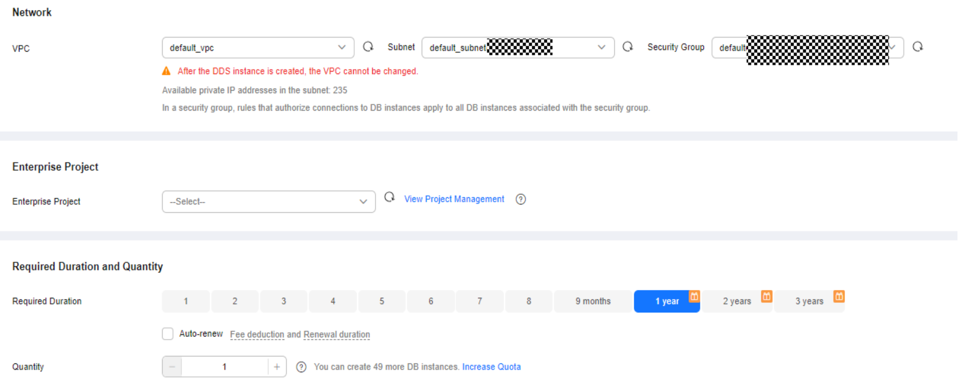

Network

VPC: default_vpc Subnet: default_subnet[redacted] Security Group: default[redacted]

After the DDS instance is created, the VPC cannot be changed.

Available private IP addresses in the subnet: 235

In a security group, rules that authorize connections to DB instances apply to all DB instances associated with the security group.

Enterprise Project

Enterprise Project: --Select-- View Project Management

Required Duration and Quantity

Required Duration: 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 months | 1 year | 2 years | 3 years (selected: 1 year)

□ Auto-renew Fee deduction and Renewal duration

Quantity: 1

You can create 49 more DB instances. Increase Quota

Table 1-9 Network settings

Parameter	Description
VPC	The VPC where your DB instances are located. A VPC isolates networks for different services. It allows you to easily manage and configure private networks and change network configurations. You need to create or select the required VPC. For details, see **Creating a VPC** in the *Virtual Private Cloud User Guide*. For details about the constraints on the use of VPCs, see **Connection Methods**. If there are no VPCs available, DDS creates one for you by default. **NOTE** After the DDS instance is created, the VPC cannot be changed.
Enterprise Project	Only enterprise users can use this function. To use this function, contact customer service. An enterprise project is a cloud resource management mode, in which cloud resources and members are centrally managed by project. Select an enterprise project from the drop-down list. The default project is **default**. For more information about enterprise project, see **Project Management** in *Enterprise Management User Guide*. To customize an enterprise project, click **Enterprise** in the upper right corner of the console. The **Enterprise Management** page is displayed. For details, see **Creating an Enterprise Project** in *Enterprise Management User Guide*.

Table 1-10 Required duration and quantity

Parameter	Description
Required Duration	The system will automatically calculate the fee based on the validity period you have selected.
Auto-renew	● By default, this option is not selected. ● If you select this option, the auto-renew cycle is determined by the length of the subscription.
Quantity	The purchase quantity depends on the replica set instance quota. If your current quota does not allow you to purchase the required number of instances, you can apply for increasing the quota as prompted. Yearly/Monthly instances that were purchased in batches have the same specifications except for the instance name and ID.

Step 3 On the displayed page, confirm the instance details.

- For yearly/monthly instances
  - If you need to modify the specifications, click **Previous** to return to the previous page.
  - If you do not need to modify the specifications, read and agree to the service agreement and click **Pay Now** to go to the payment page and complete the payment.
- For pay-per-use instances
  - If you need to modify the specifications, click **Previous** to return to the previous page.
  - If you do not need to modify the specifications, read and agree to the service agreement and click **Submit** to start creating the instance.

Step 4 Click **Back to Instance List**. After a DDS instance is created, you can view and manage it on the **Instances** page.

- When an instance is being created, the status displayed in the **Status** column is **Creating**. This process takes about 15 minutes. After the creation is complete, the status changes to **Available**.
- DDS enables the automated backup policy by default. After an instance is created, you can modify or disable the automated backup policy. An automated full backup is immediately triggered after the creation of an instance.

----End

1.2.2 Custom Config

This section describes how to purchase a replica set instance in custom mode on the management console. You can customize the computing resources and storage space of a replica set instance based on your service requirements. In addition, you can configure advanced settings, such as slow query log and automated backup.

Precautions

Each account can create up to 50 replica set instances.

Prerequisites

- You have **registered a Huawei ID and enabled Huawei Cloud services**.
- To display whether the disk is encrypted in the DB instance list, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket**.
- If you want compute and network resources dedicated to your exclusive use, **enable a DeC** and **apply for DCC resources**. Then, you can create DDS instances. Click 📍 in the upper left corner and select a region and a project.

NOTE

You will be additionally charged for using DeC. Only pay-per-use replica set instances can be purchased through DeC.

Procedure

Step 1 Go to the **Custom Config** page.

Step 2 On the displayed page, select a billing mode and configure information about your DB instance. Then, click **Next**.

Figure 1-9 Basic configurations

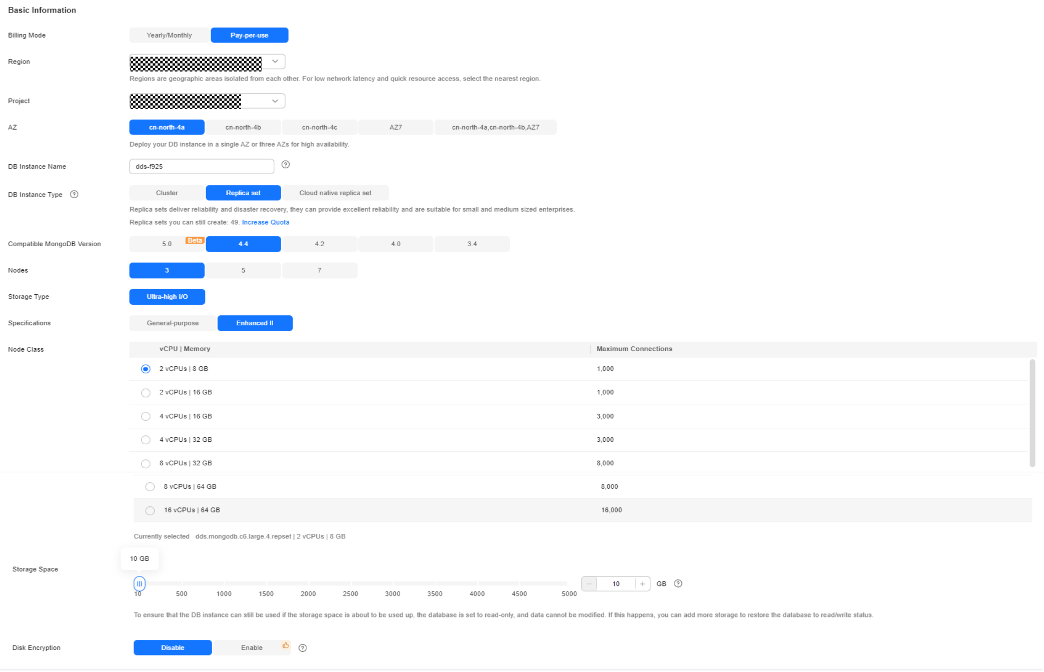

Basic Information

Billing Mode: Pay-per-use

Region: [redacted]

Regions are geographic areas isolated from each other. For low network latency and quick resource access, select the nearest region.

Project: [redacted]

AZ: cn-north-4a (cn-north-4b, cn-north-4c, AZ7, cn-north-4a,cn-north-4b,AZ7)

Deploy your DB instance in a single AZ or three AZs for high availability.

DB Instance Name: dds-f925

DB Instance Type: Replica set

Replica sets deliver reliability and disaster recovery, they can provide excellent reliability and are suitable for small and medium sized enterprises.

Replica sets you can still create: 49. Increase Quota

Compatible MongoDB Version: 4.4

Nodes: 3

Storage Type: Ultra-high I/O

Specifications: Enhanced II

Node Class

	vCPU | Memory	Maximum Connections
●	2 vCPUs | 8 GB	1,000
○	2 vCPUs | 16 GB	1,000
○	4 vCPUs | 16 GB	3,000
○	4 vCPUs | 32 GB	3,000
○	8 vCPUs | 32 GB	8,000
○	8 vCPUs | 64 GB	8,000
○	16 vCPUs | 64 GB	16,000

Currently selected dds.mongodb.c6.large.4.repset | 2 vCPUs | 8 GB

Storage Space: 10 GB

10 500 1000 1500 2000 2500 3000 3500 4000 4500 5000

10 GB

To ensure that the DB instance can still be used if the storage space is about to be used up, the database is set to read-only, and data cannot be modified. If this happens, you can add more storage to restore the database to read/write status.

Disk Encryption: Disable

Table 1-11 Billing Mode

Parameter	Description
Billing Mode	Select a billing mode, **Yearly/Monthly** or **Pay-per-use**. ● For yearly/monthly instances – Specify **Required Duration**, and the system deducts the fees incurred from your account based on the service price. – If you do not expect to continue using the instance much after it expires, you can change the billing mode from yearly/monthly to pay-per-use. For details, see **Yearly/Monthly to Pay-per-Use**. **NOTE** Instances billed on a yearly/monthly basis cannot be deleted. They can only be unsubscribed from. For details, see **Unsubscribing from a Yearly/Monthly Instance**. ● For pay-per-use instances – You are billed for usage based on how much time the service is in use. – If you expect to use the service extensively over a long period of time, you can change its billing mode from pay-per-use to yearly/monthly to reduce costs. For details, see **Pay-per-Use to Yearly/Monthly**.
Region	The region where the resource is located. **NOTE** Instances deployed in different regions cannot communicate with each other through a private network, and you cannot change the region of an instance once it is purchased. Exercise caution when selecting a region.
Project	The project corresponds to the current region and can be changed.

Parameter	Description
AZ	An AZ is a part of a region with its own independent power supply and network. AZs are physically isolated but can communicate through internal network connections. Instances can be deployed in a single AZ or three AZs. ● If your service requires low network latency between instances, you deploy the components of the instance in the same AZ. If you select a single AZ to deploy your instance, anti-affinity deployment is used by default. With an anti-affinity deployment, your primary, secondary, and hidden nodes are deployed on different physical machines for high availability. ● If you want to deploy an instance across AZs for disaster recovery, select three AZs. In this deployment mode, the primary, secondary, and hidden nodes are evenly distributed across three AZs. **NOTE** The 3-AZ deployment is not available in all regions. If the 3-AZ option is not displayed on the page for you to buy an instance, try a different region.
DB Instance Name	● The instance name that you specify after the purchase. The instance name must contain 4 to 64 characters and must start with a letter. It is case sensitive and can contain letters, digits, hyphens (-), and underscores (_). It cannot contain other special characters. ● The instance name can be the same as an existing instance name. ● If you buy a batch of instances at once, a 4-digit numerical suffix will be added to the instance names, starting with **-0001**. If you later make another batch purchase, the new instance names will be numbered first using any suffixes missing from the sequence of your existing instances, and then continuing on from where your last batch purchase left off. For example, a batch of 3 instances gets the suffixes **-0001**, **-0002**, and **-0003**. If you deleted instance **0002** and then bought 3 more instances, the new instances would get the suffixes **-0002**, **-0004**, and **-0005**. ● After the DB instance is created, you can change its name. For details, see **Changing an Instance Name**.
DB Instance Type	Select **Replica set**. A replica set consists of the primary node, secondary node, and hidden node. If a primary node goes down or becomes faulty, a secondary node is automatically assigned to the primary role and continues normal operation. If a secondary node is unavailable, a hidden node will take the role of the secondary to ensure high availability.

Parameter	Description
Primary AZ	Select the AZ housing the primary/standby role. **NOTE** This parameter is available when **AZ** is set to multiple AZs.
Standby AZ	Select the AZ housing the primary/standby role. **NOTE** This parameter is available when **AZ** is set to multiple AZs.
Compatible MongoDB Version	● 5.0 ● 4.4 ● 4.2 ● 4.0 ● 3.4
Nodes	You can create a three-node, five-node, or seven-node replica set instance.
CPU Type	DDS supports x86 and Kunpeng CPU architectures. **NOTE** This parameter is available only for MongoDB 4.0 and 3.4. You do not need to set this parameter for other versions. The default value is **x86**. ● x86 x86 CPUs use the Complex Instruction Set Computing (CISC) instruction set. Each instruction can be used to execute low-level hardware operations. CISC instructions vary in length, and tend to be complicated and slow compared to Reduced Instruction Set Computing (RISC). ● Kunpeng The Kunpeng CPU architecture uses RISC. The RISC instruction set is smaller and faster than CISC, thanks to the simplified architecture. Kunpeng CPUs also offer a better balance between power and performance than x86. Kunpeng CPUs offer a high density, low power option that is more cost effective for heavy workloads.
Storage Type	The storage type can be **Ultra-high I/O** and **Extreme SSD** for non-DeC users. For DeC users, the supported storage types depend on the selected resource type. ● If you select **EVS** for **Resource Type**, **Storage Type** is set to **Cloud SSD**. ● If you select **DSS** for **Resource Type**, **Storage Type** can be set to **Common I/O**, **High I/O**, or **Cloud SSD**.

Parameter	Description
Storage Engine	- WiredTiger WiredTiger is the default storage engine of DDS 3.4 and 4.0. WiredTiger provides different granularity concurrency control and compression mechanism for data management. It can provide the best performance and storage efficiency for different kinds of applications. - RocksDB RocksDB is the default storage engine of DDS 4.2. RocksDB supports efficient point lookup, range scan, and high-speed write. RocksDB can be used as the underlying data storage engine of MongoDB and is suitable for scenarios with a large number of write operations.
Specifications	With an x86 architecture, you have the following options: - General-purpose (s6): S6 instances are suitable for applications that require moderate performance generally but occasional bursts of high performance, such as light-workload web servers, enterprise R&D and testing environments, and low- and medium-performance databases. - Enhanced II (c6): C6 instances have multiple technologies optimized to provide stable powerful compute performance. 25 GE intelligent high-speed NICs are used to provide ultra-high bandwidth and throughput, making it an excellent choice for heavy-load scenarios. It is suitable for websites, web applications, general databases, and cache servers that have higher performance requirements for compute and network resources; and medium- and heavy-load enterprise applications.
Node Class	For details about the instance specifications, see **Instance Specifications**. For details about the performance data of DB instances of different specifications, see **Performance White Paper**. If the CPU or memory of a created DB instance cannot meet service requirements, you can change it on the management console. For details, see **Changing a Replica Set Instance Class**.
Storage Space	The storage space ranges from 10 GB to 5,000 GB. The value must be an integer multiple of 10. You can scale up an instance after it is created. For details, see **Scaling Up a Replica Set Instance**. **NOTE** - If the storage space you purchased exceeds 600 GB and the remaining storage space is 18 GB, the instance becomes **Read-only**. - If the storage space you purchased is less than 600 GB and the storage space usage reaches 97%, the instance becomes **Read-only**. In these cases, delete unnecessary resources or expand the capacity.

Parameter	Description
Disk Encryption	● **Disabled**: Disable encryption. ● **Enabled**: Enable encryption. This feature improves data security but slightly affects read/write performance. **Key Name**: Select or create a private key, which is the tenant key. **NOTE** ● After an instance is created, the disk encryption status and the key cannot be changed. Disk encryption will not encrypt backup data stored in OBS. To enable backup data encryption, contact customer service. ● To check whether the disk is encrypted, you can view **Disk Encrypted** in the DB instance list. ● If disk encryption or backup data encryption is enabled, keep the key properly. Once the key is disabled, deleted, or frozen, the database will be unavailable and data may not be restored. If disk encryption is enabled but backup data encryption is not enabled, you can **restore data to a new instance from backups**. If both disk encryption and backup data encryption are enabled, data cannot be restored. ● For details about how to create a key, see "**Creating a CMK**" in *Data Encryption Workshop User Guide*. ● Disk encryption supports only the AES_256 key algorithm.

Figure 1-10 Administrator settings

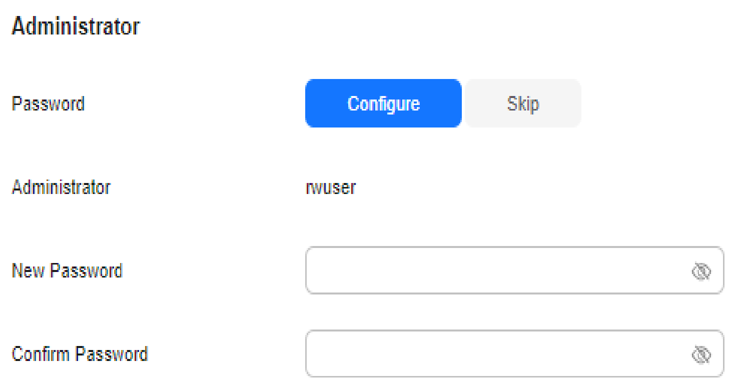

Administrator

Password: Configure | Skip

Administrator: rwuser

New Password:

Confirm Password:

Table 1-12 Administrator settings

Parameter	Description
Password	● Configure Enter and confirm the new administrator password. After an instance is created, you can connect to the instance using the password. ● Skip To log in, you will have to reset the password later on the **Basic Information** page. If you need to connect to an instance after it is created, locate the instance and choose **More** > **Reset Password** in the **Operation** column to set a password for the instance first.
Administrator	The default account is **rwuser**.
Administrator Password	Set a password for the administrator. The password must be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and at least one of the following special characters: ~!@#%^*-_=+?()$ Keep this password secure. If lost, the system cannot retrieve it for you.
Confirm Password	Enter the administrator password again.

Figure 1-11 Network, Required Duration, and Quantity

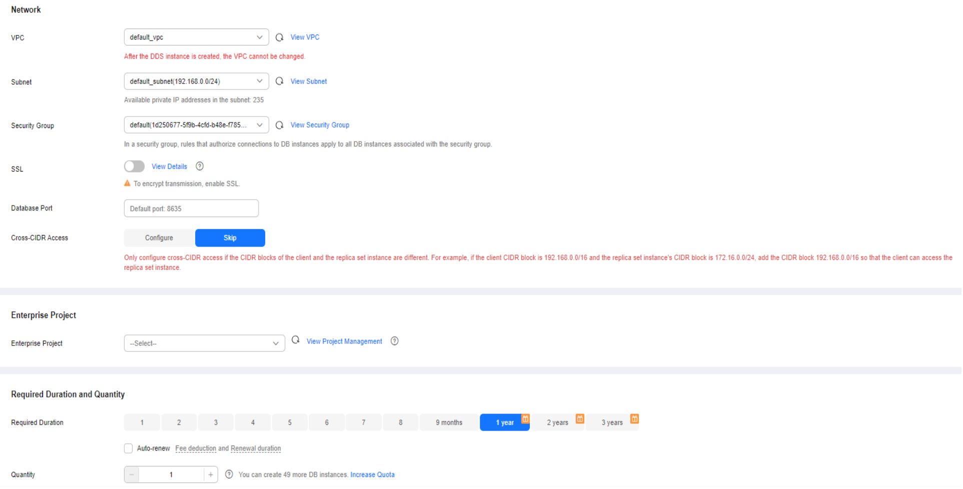

Network

VPC: default_vpc View VPC

After the DDS instance is created, the VPC cannot be changed.

Subnet: default_subnet(192.168.0.0/24) View Subnet

Available private IP addresses in the subnet: 235

Security Group: default(1d250677-5f9b-4cfd-b48e-f785... View Security Group

In a security group, rules that authorize connections to DB instances apply to all DB instances associated with the security group.

SSL: View Details

To encrypt transmission, enable SSL.

Database Port: Default port: 8635

Cross-CIDR Access: Configure [✓ Skip]

Only configure cross-CIDR access if 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.

Enterprise Project

Enterprise Project: --Select-- View Project Management

Required Duration and Quantity

Required Duration: 1 2 3 4 5 6 7 8 9 months [✓ 1 year] 2 years 3 years

□ Auto-renew Fee deduction and Renewal duration

Quantity: − 1 +

You can create 49 more DB instances. Increase Quota

Table 1-13 Network

Parameter	Description
VPC	The VPC where your DB instances are located. A VPC isolates networks for different services. It allows you to easily manage and configure private networks and change network configurations. You will need to create or select the required VPC. For details about how to create a VPC, see "Creating a VPC" in *Virtual Private Cloud User Guide*. For details about the constraints on the use of VPCs, see **Connection Methods**. If there are no VPCs available, DDS creates one for you by default. VPC owners can share the subnets in a VPC with one or multiple accounts through Resource Access Manager (RAM). This allows for more efficient use of network resources and reduces O&M costs. For more information about VPC subnet sharing, see **VPC Sharing** in the *Virtual Private Cloud User Guide*. **NOTE** After the DDS instance is created, the VPC cannot be changed.
Subnet	A subnet provides dedicated network resources that are logically isolated from other networks for security reasons. After the instance is created, you can change the private IP address assigned by the subnet. For details, see **Changing a Private IP Address**. **NOTE** IPv6 subnets are not supported. You are advised to create and select IPv4 subnets.
Security Group	A security group controls access between DDS and other services. If there are no security groups available, DDS creates one for you by default. **NOTE** ● Ensure that there is a security group rule configured that allows clients to access instances. For example, select an inbound TCP rule with the default port 8635, and enter a subnet IP address or select a security group that the instance belongs to. ● When creating a DB instance, you can select multiple security groups. For better network performance, you are advised to select no more than five security groups. In such a case, the access rules of all the selected security groups apply on the instance.
SSL	Secure Sockets Layer (SSL) encrypts connections between clients and servers, preventing data from being tampered with or stolen during transmission. You can enable SSL to improve data security. After an instance is created, you can connect to it using SSL.

Parameter	Description
Database Port	The default DDS port is 8635, but this port can be modified if necessary. If you change the port, add a corresponding security group rule to allow access to the instance. **NOTE** ● For details about how to change a database port, see **Changing a Database Port**.
Cross-CIDR Access	● Configure If a client and a replica set instance are deployed in different CIDR blocks and the client is not in 192.168.0.0/16, 172.16.0.0/24, or 10.0.0.0/8, configure Cross-CIDR Access for the instance to communicate with the client. **NOTE** – To ensure the ECS and the DB instance can communicate with each other, configure the connection by referring to **VPC Peering Connection Overview**. – Up to 30 CIDR blocks can be configured, and each of them can overlap but they cannot be the same. That is, the source CIDR blocks can overlap but cannot be the same. The CIDR blocks cannot start with 127. The allowed IP mask ranges from 8 to 32. ● Skip Configure the CIDR block of the client later. After a DB instance is created, you can configure cross-CIDR access by referring to **Configuring Cross-CIDR Access**.
Enterprise Project	Only enterprise users can use this function. To use this function, contact customer service. An enterprise project is a cloud resource management mode, in which cloud resources and members are centrally managed by project. Select an enterprise project from the drop-down list. The default project is **default**.

Figure 1-12 Advanced settings

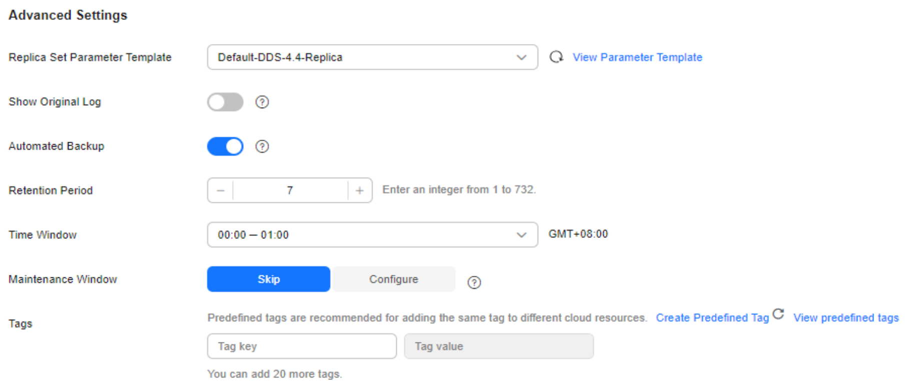

Advanced Settings

Replica Set Parameter Template: Default-DDS-4.4-Replica

Show Original Log: off

Automated Backup: on

Retention Period: 7 (Enter an integer from 1 to 732.)

Time Window: 00:00 — 01:00 GMT+08:00

Maintenance Window: Skip

Tags

Predefined tags are recommended for adding the same tag to different cloud resources.

You can add 20 more tags.

Table 1-14 Advanced settings

Parameter	Description
Replica Set Parameter Template	The parameters that apply to the replica set instances. After an instance is created, you can change the parameter template you configured for the instance to bring out the best performance. For details, see **Editing a Parameter Template**.
Show Original Log	If Show Original Log is enabled, the original slow query logs will be displayed, and the logs will be transferred to an OBS bucket. By default, the system automatically deletes logs from the OBS bucket after 30 days, and the retention period cannot be changed.
Automated Backup	DDS enables an automated backup policy by default, but you can disable it after an instance is created. An automated full backup is immediately triggered after the creation of an instance. For details, see **Configuring an Automated Backup Policy**.
Retention Period (days)	**Retention Period** refers to the number of days that data is kept. You can increase the retention period to improve data reliability. The backup retention period is from 1 to 732 days.
Time Window	The backup interval is 1 hour.

Parameter	Description
Tags	(Optional) You can add tags to DDS instances so that you can quickly search for and filter specified instances by tag. Each DDS instance can have up to 20 tags. If your organization has configured tag policies for DDS, add tags to DB instances based on the policies. If a tag does not comply with the policies, DB instance creation may fail. Contact your organization administrator to learn more about tag policies. ● Create a tag. You can create tags on the DDS console. A tag key and a value are required when you create a tag. Key: This parameter is mandatory. – Each tag key must be unique for each instance. – A tag key consists of up to 36 characters. – The key must consist of only digits, letters, underscores (_), and hyphens (-). Value: This parameter is optional. – The value consists of up to 43 characters. – The value must consist of only digits, letters, underscores (_), periods (.), and hyphens (-). ● Add a predefined tag. Predefined tags can be used to identify multiple cloud resources. To tag a cloud resource, you can select a created predefined tag from the drop-down list, without entering a key and value for the tag. For example, if a predefined tag has been created, its key is Usage and value is Project1. When you configure the key and value for a cloud resource, the created predefined tag will be displayed on the page. After an instance is created, you can click the instance name to view its tags. On the **Tags** page, you can also **modify or delete the tags**. In addition, you can quickly **search for and filter specified instances by tag**. You can add a tag to an instance after the instance is created. For details, see **Adding a Tag**.

If you have any question about the price, click **Price Details**.

NOTE

Instance performance depends on the specifications you select during creation. The hardware configuration items that can be selected include the instance class and storage space.

Step 3 On the displayed page, confirm the instance details.

- For yearly/monthly instances
  - If you need to modify the specifications, click **Previous** to return to the previous page.
  - If you do not need to modify the specifications, read and agree to the service agreement and click **Pay Now** to go to the payment page and complete the payment.
- For pay-per-use instances
  - If you need to modify the specifications, click **Previous** to return to the previous page.
  - If you do not need to modify the specifications, read and agree to the service agreement and click **Submit** to start creating the instance.

Step 4 Click **Back to Instance List**. After a DDS instance is created, you can view and manage it on the **Instances** page.

- When an instance is being created, the status displayed in the **Status** column is **Creating**. This process takes about 15 minutes. After the creation is complete, the status changes to **Available**.
- Yearly/Monthly instances that were purchased in batches have the same specifications except for the instance name and ID.

----End

2 Connecting to a DB Instance

2.1 Connecting to a Cluster Instance

2.1.1 Connection Methods

You can access DDS over private or public networks.

Table 2-1 Connection methods

Method	IP Address	Scenario	Description
DAS	Not required	DAS provides a GUI and allows you to perform visualized operations on the console. SQL execution, advanced database management, and intelligent O&M are all available to make database management simple, secure, and intelligent. By default, the permission to connect to DAS is enabled.	● Easy to use, secure, advanced, and intelligent ● Recommended

Method	IP Address	Scenario	Description
Private network	Private IP address	DDS provides a private IP address by default. If your applications are running on an ECS in the same region and VPC as your DDS instance, you are advised to use a private IP address to connect the ECS to your DDS instances.	● Secure and excellent performance ● For faster transmission and improved security, you are advised to migrate your applications to an ECS that is in the same subnet as your DDS instance and use a private IP address to access the instance.
Public network	EIP	● If your applications are running on an ECS that is in a different region from the one where the DDS instance is located, use an EIP to connect the ECS to your DDS instances. ● If you use a third-party device or your local device to connect to a DDS instance, you can use an EIP to connect to the DB instance.	● Low security

2.1.2 (Recommended) Connecting to Cluster Instances Through DAS

2.1.2.1 Overview

DAS provides a GUI and allows you to perform visualized operations on the console. SQL execution, advanced database management, and intelligent O&M are all available to make database management simple, secure, and intelligent. You are advised to use DAS to connect to instances.

This section describes how to buy a cluster instance on the management console and how to connect to the cluster instance through DAS.

Process

To purchase and connect to a cluster instance, perform the following steps:

1. **Buy a cluster instance.**
2. **Connect to the cluster instance through DAS.**

2.1.2.2 Connecting to a Cluster Instance Through DAS

Data Admin Service (DAS) enables you to manage DB instances on a web-based console, simplifying database management and improving working efficiency. You can connect and manage instances through DAS. By default, you have the permission required for remote login. It is recommended that you use the DAS service to connect to DB instances. DAS is secure and convenient.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

If you want compute and network resources dedicated to your exclusive use, **enable a DeC** and **apply for DCC resources**. After enabling a DeC, you can select the DeC region and project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click **Log In** in the **Operation** column.

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

Step 5 In the **Instance Login** dialog box, enter the correct information and click **Log In** to access and manage your database.

Step 6 After the login is successful, you can perform operations such as creating a database, managing accounts, and managing databases.

For details, see **Database Management**.

----End

2.1.3 Connecting to a Cluster Instance over a Private Network

2.1.3.1 Configuring Security Group Rules

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

To ensure database security and reliability, you need to configure security group rules to allow specific IP addresses and ports to access DDS instances.

You can connect to an instance by configuring security group rules in following two ways:

- If the ECS and instance are in the same security group, they can communicate with each other by default. No security group rule needs to be configured. Go to **Connecting to a Cluster Instance Using Mongo Shell (Private Network)**.

Figure 2-1 Same security group

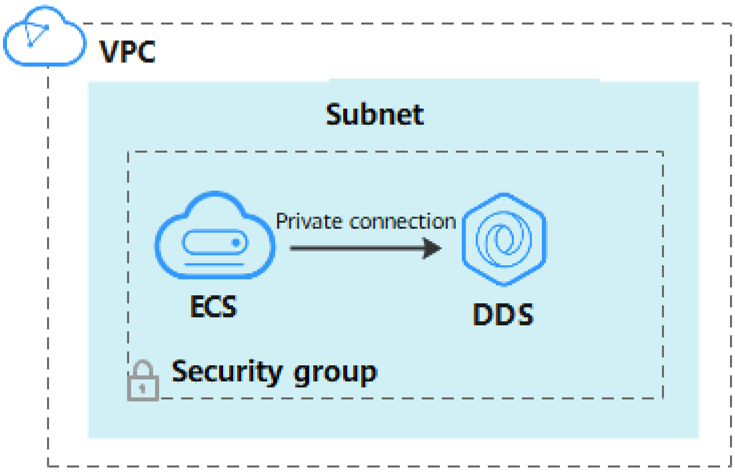

- If the ECS and instance are in different security groups, you need to configure security group rules for them, separately.

Figure 2-2 Different security groups

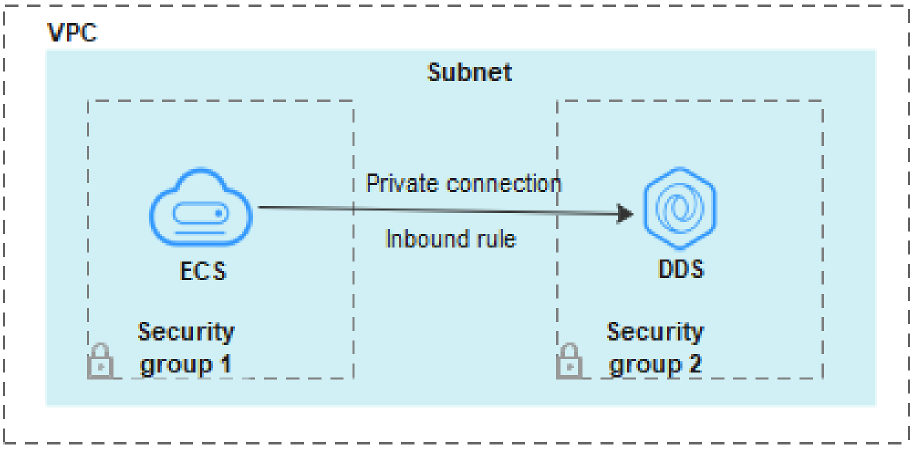

  - Instance: Configure an **inbound rule** for the security group associated with the instance.
  - ECS: The default security group rule allows all outbound data packets. In this case, you do not need to configure a security group rule for the ECS. If not all traffic is allowed to reach the instance, configure an **outbound** rule for the ECS.

This section describes how to configure an **inbound** rule for an instance.

Precautions

- By default, an account can create up to 500 security group rules.
- Too many security group rules will increase the first packet latency, so a maximum of 50 rules for each security group is recommended.
- By default, one DDS instance is associated with only one security group.
- DDS allows you to associate multiple security groups to a DB instance. You can apply for the service based on your service requirements. For better network performance, you are advised to select no more than five security groups.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Network Information** area on the **Basic Information** page, click the security group.

Figure 2-3 Security Group

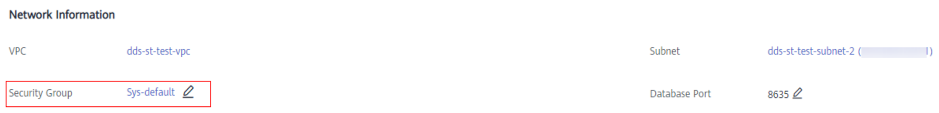

Network Information

VPC	dds-st-test-vpc	Subnet	dds-st-test-subnet-2 ([redacted])
Security Group	Sys-default	Database Port	8635

You can also choose **Connections** in the navigation pane on the left. On the **Private Connection** tab, in the **Security Group** area, click the security group name.

Figure 2-4 Security Group

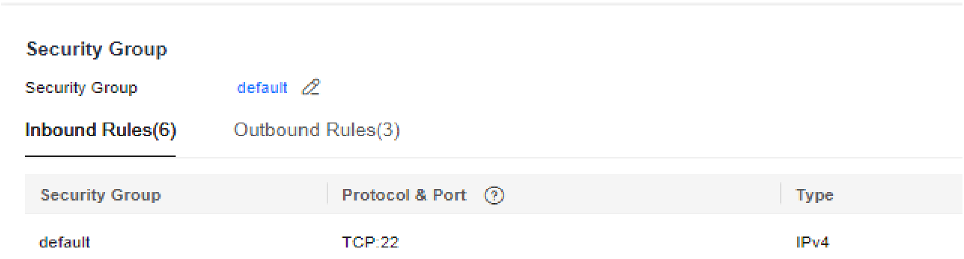

Security Group

Security Group default

Inbound Rules(6) Outbound Rules(3)

Security Group	Protocol & Port	Type
default	TCP:22	IPv4

Step 6 On the **Security Group** page, locate the target security group and click **Manage Rule** in the **Operation** column.

Step 7 On the **Inbound Rules** tab, click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Step 8 Add a security group rule as prompted.

Figure 2-5 Add Inbound Rule

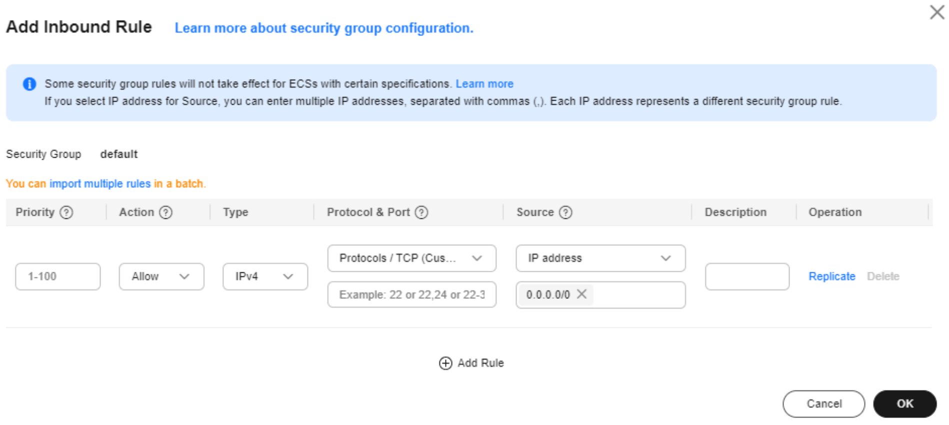

Add Inbound Rule Learn more about security group configuration.

Some security group rules will not take effect for ECSs with certain specifications. Learn more
If you select IP address for Source, you can enter multiple IP addresses, separated with commas (,). Each IP address represents a different security group rule.

Security Group default

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Source	Description	Operation
1-100	Allow	IPv4	Protocols / TCP (Cus... Example: 22 or 22,24 or 22-3	IP address 0.0.0.0/0		Replicate Delete

Table 2-2 Inbound rule settings

Parameter	Description	Example
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default priority is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	The security group rule actions. A rule with a deny action overrides another with an allow action if the two rules have the same priority.	Allow
Protocol & Port	The network protocol required for access. Available options: **TCP**, **UDP**, **ICMP**, or **GRE**	TCP
	Port: the port on which you wish to allow access to DDS. The default port is 8635. The port ranges from 2100 to 9500 or can be 27017, 27018, or 27019.	8635
Type	IP address type. Only **IPv4** and **IPv6** are supported.	IPv4

Parameter	Description	Example
Source	Specifies the supported IP address, security group, and IP address group, which allow access from IP addresses or instances in other security group. Example: - Single IP address: 192.168.10.10/32 - IP address segment: 192.168.1.0/24 - All IP addresses: 0.0.0.0/0 - Security group: sg-abc - IP address group: ipGroup-test If you enter a security group, all ECSs associated with the security group comply with the created rule. For more information about IP address groups, see **IP Address Group Overview**.	0.0.0.0/0
Description	(Optional) Provides supplementary information about the security group rule. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

Step 9 Click **OK**.

----End

2.1.3.2 Connecting to a Cluster Instance Using Mongo Shell (Private Network)

Mongo shell is the default client for the MongoDB database server. You can use Mongo Shell to connect to DB instances, and query, update, and manage data in databases. DDS is compatible with MongoDB. Mongo Shell is a part of the MongoDB client. To use Mongo Shell, download and install the MongoDB client first, and then use the Mongo shell to connect to the DB instance.

By default, a DDS instance provides a private IP address. If your applications are deployed on an ECS and are in the same region and VPC as DDS instances, you can connect to DDS instances using a private IP address to achieve a fast transmission rate and high security.

This section describes how to use Mongo Shell to connect to a cluster instance over a private network.

You can connect to an instance using an SSL connection or an unencrypted connection. The SSL connection is encrypted and more secure. To improve data transmission security, connect to instances using SSL.

Prerequisites

1. For details about how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS**.
2. You have installed the MongoDB client on the ECS. To ensure successful authentication, install the MongoDB client of the same version as the target instance.

   For details about how to install a MongoDB client, see **How Can I Install a MongoDB Client?**
3. The ECS can communicate with the DDS instance. For details, see **Configuring Security Group Rules**.

SSL Connection

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Connections**.

Step 6 In the **Basic Information** area, click 📥 next to the **SSL** field.

Step 7 Upload the root certificate to the ECS to be connected to the instance.

The following describes how to upload the certificate to a Linux and Windows ECS:

- In Linux, run the following command:

  ```
  scp <IDENTITY_FILE><REMOTE_USER>@<REMOTE_ADDRESS>:<REMOTE_DIR>
  ```

  NOTE

  - **IDENTITY_FILE** is the directory where the root certificate resides. The file access permission is 600.
  - **REMOTE_USER** is the ECS OS user.
  - **REMOTE_ADDRESS** is the ECS address.
  - **REMOTE_DIR** is the directory of the ECS to which the root certificate is uploaded.
- In Windows, upload the root certificate using a remote connection tool.

Step 8 Connect to the instance in the directory where the MongoDB client is located.

Method 1: Using the private HA connection address (recommended)

DDS provides a private HA connection address that consists of IP addresses and ports of all dds mongos nodes in a cluster instance. You can use this address to connect to the cluster instance to improve availability of the cluster instance.

Example command:

./mongo *<Private HA connection address>* **--ssl --sslCAFile** *<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **Private HA Connection Address**: On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. Click the **Private Connection** tab and obtain the connection address of the current instance from the **Private HA Connection Address** field.

  Figure 2-6 Obtaining the private HA connection address

  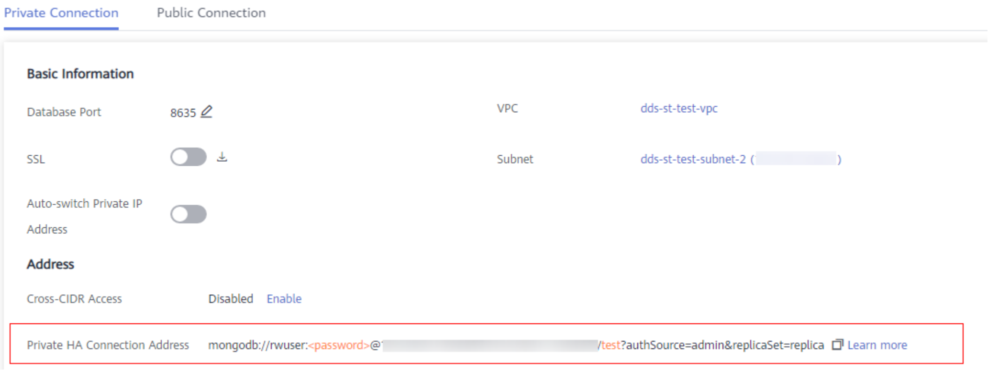

  Basic Information

  Database Port 8635

  VPC dds-st-test-vpc

  SSL

  Subnet dds-st-test-subnet-2 ([redacted])

  Auto-switch Private IP Address

  Address

  Cross-CIDR Access Disabled Enable

  Private HA Connection Address mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

  The format of the private HA connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

  mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin**

  The following table lists the required parameters in the private HA address.

  Table 2-3 Parameter information

Parameter	Description
rwuser	Database username
<password>	Password for the database username. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.

Parameter	Description
192.168.***.***:8635,192.168.***.***:8635	IP addresses and ports of the dds mongos nodes of the cluster instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: To ensure that the internal communication of the cluster does not occupy resources such as the user IP address and bandwidth, the cluster certificate is generated using the internal management IP address. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

Command example:

./mongo mongodb://rwuser: *<password>@192.168.xx.xx:8635,192.168.xx.xx:8635/* **test?authSource=admin --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

Method 2: Using the private HA connection address (user-defined database and account)

Example command:

./mongo *<Private HA connection address>* **--ssl --sslCAFile** *<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **Private HA Connection Address**: On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. Click the **Private Connection** tab and obtain the connection address of the current instance from the **Private HA Connection Address** field.

Figure 2-7 Obtaining the private HA connection address

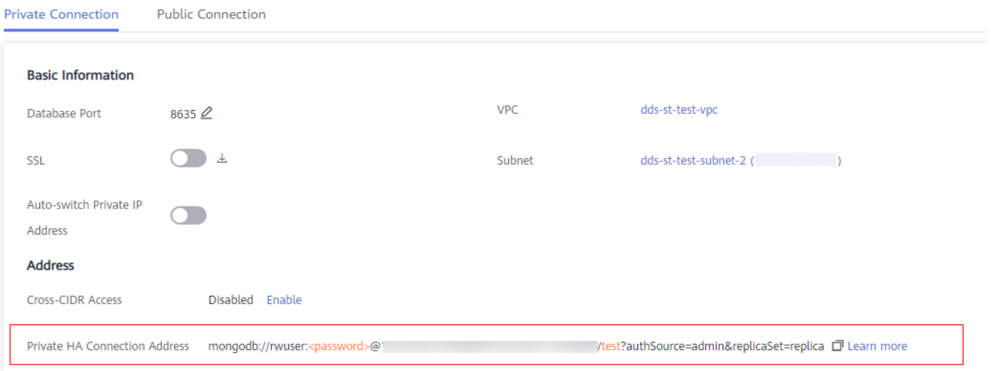

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet-2 ([redacted])

Auto-switch Private IP Address

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

The format of the obtained private HA connection address is as follows:

mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin**

The following table lists the required parameters in the private HA address.

Table 2-4 Parameter information

Parameter	Description
rwuser	Database username. The default value is **rwuser**. You can change the value to the username based on your service requirements.
<password>	Password for the database username. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.***.***:8635,192.168.***.***:8635	IP addresses and ports of the dds mongos nodes of the cluster instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** is **admin**. **NOTE** If you use a user-defined database for authentication, change the authentication database in the HA connection address to the name of the user-defined database. In addition, replace **rwuser** with the username created in the user-defined database.

- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: To ensure that the internal communication of the cluster does not occupy resources such as the user IP address and bandwidth, the cluster certificate is generated using the internal management IP address. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

For example, if you create a user-defined database **Database** and user **test1** in the database, the connection command is as follows:

./mongo mongodb://test1:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/Database?authSource=Database --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

Method 3: Using a private IP address

Example command:

./mongo --host *<DB_HOST>* **--port** *<DB_PORT>* **-u** *<DB_USER>* **-p --authenticationDatabase admin --ssl --sslCAFile** *<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **DB_HOST** is the IP address of the dds mongos node of the cluster instance to be connected.

  Click the instance name. On the **Basic Information** page, choose **Connections** > **Private Connection**, obtain the private IP address of the dds mongos node on the **dds mongos** tab in the **Node Information** area.

  Figure 2-8 Obtaining the private IP address

  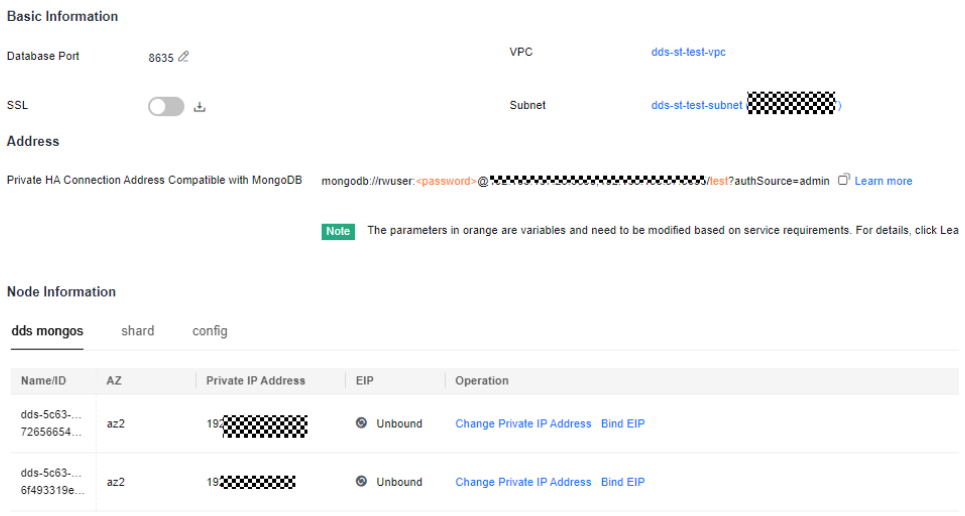

  Basic Information

  Database Port 8635

  VPC dds-st-test-vpc

  SSL

  Subnet dds-st-test-subnet ([redacted])

  Address

  Private HA Connection Address Compatible with MongoDB mongodb://rwuser:<password>@[redacted]/test?authSource=admin Learn more

  Note The parameters in orange are variables and need to be modified based on service requirements. For details, click Lea

  Node Information

Name/ID	AZ	Private IP Address	EIP	Operation
dds-5c63-... 72656654...	az2	[redacted]	Unbound	Change Private IP Address Bind EIP
dds-5c63-... 6f493319e...	az2	[redacted]	Unbound	Change Private IP Address Bind EIP

- **DB_PORT** is the port of the instance to be connected. The default port is 8635.

  Click the instance name. On the **Basic Information** page, choose **Connections**. On the **Private Connection** tab, obtain the database port information in the **Database Port** field in the **Basic Information** area.

  Figure 2-9 Obtaining the port

  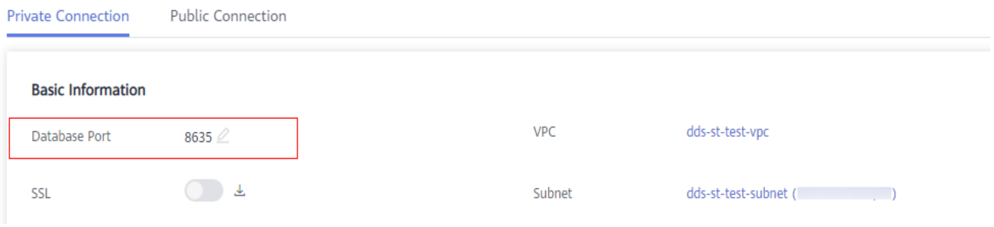

  Basic Information

  Database Port 8635

  VPC dds-st-test-vpc

  SSL

  Subnet dds-st-test-subnet ([redacted])

- **DB_USER** is the database user. The default value is **rwuser**.
- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: To ensure that the internal communication of the cluster does not occupy resources such as the user IP address and bandwidth, the cluster certificate is generated using the internal management IP address. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

Enter the database account password when prompted:

```
Enter password:
```

Command example:

./mongo --host 192.168.1.6 --port 8635 -u rwuser -p --authenticationDatabase admin --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames

Step 9 Check the connection result. If the following information is displayed, the connection is successful.

```
mongos>
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Connect to the ECS.

Step 2 Connect to the instance in the directory where the MongoDB client is located.

Method 1: Private HA connection address (recommended)

Example command:

./mongo "*<Private HA Connection Address>***"**

Private HA Connection Address: On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. Click the **Private Connection** tab and obtain the connection address of the current instance from the **Private HA Connection Address** field.

Figure 2-10 Obtaining the private HA connection address

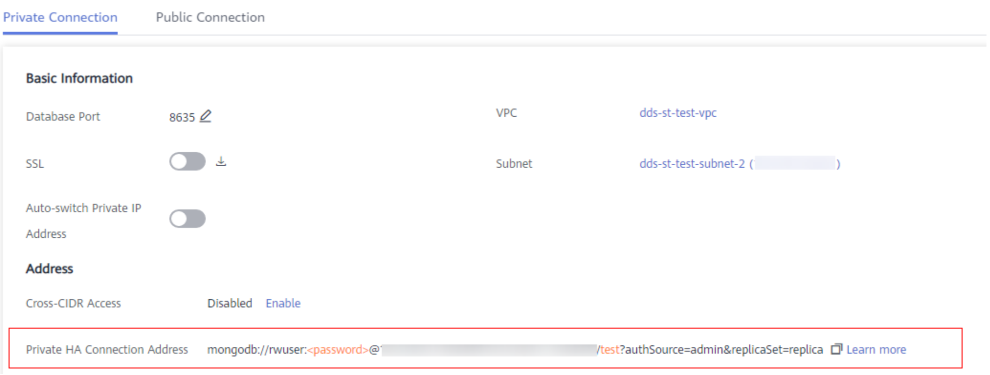

Basic Information

Database Port 8635

SSL

Auto-switch Private IP Address

VPC dds-st-test-vpc

Subnet dds-st-test-subnet-2 ([redacted])

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

The format of the private HA connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin**

The following table lists the required parameters in the private HA address.

Table 2-5 Parameter information

Parameter	Description
rwuser	Database username
<password>	Password for the database username. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.***.***:8635,192.168.***.***:8635	IP addresses and ports of the dds mongos nodes of the cluster instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

Command example:

./mongo mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin**

Method 2: Private HA connection (user-defined database and account)

Example command:

./mongo "*<Private HA Connection Address>***"**

Private HA Connection Address: On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. Click the **Private Connection** tab and obtain the connection address of the current instance from the **Private HA Connection Address** field.

Figure 2-11 Obtaining the private HA connection address

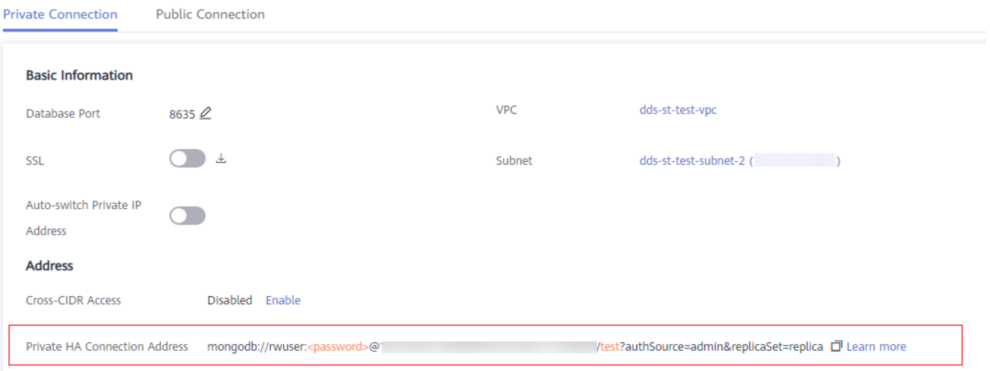

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet-2 ([redacted])

Auto-switch Private IP Address

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica

The format of the obtained private HA connection address is as follows:

mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin**

The following table lists the required parameters in the private HA address.

Table 2-6 Parameter information

Parameter	Description
rwuser	Database username. The default value is **rwuser**. You can change the value to the username based on your service requirements.
<password>	Password for the database username. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.***.***:8635,192.168.***.***:8635	IP addresses and ports of the dds mongos nodes of the cluster instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** is **admin**. **NOTE** If you use a user-defined database for authentication, change the authentication database in the HA connection address to the name of the user-defined database. In addition, replace **rwuser** with the username created in the user-defined database.

For example, if you create a user-defined database **Database** and user **test1** in the database, the connection command is as follows:

./mongo mongodb://test1:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/Database?authSource=Database**

Method 3: Using a private IP address

Example command:

./mongo --host *<DB_HOST>* **--port** *<DB_PORT>* **-u** *<DB_USER>* **-p --authenticationDatabase admin**

Parameter description:

- **DB_HOST** is the IP address of the dds mongos node of the cluster instance to be connected.

  Click the instance name. On the **Basic Information** page, choose **Connections** > **Private Connection**, obtain the private IP address of the dds mongos node on the **dds mongos** tab in the **Node Information** area.

  Figure 2-12 Obtaining the private IP address

  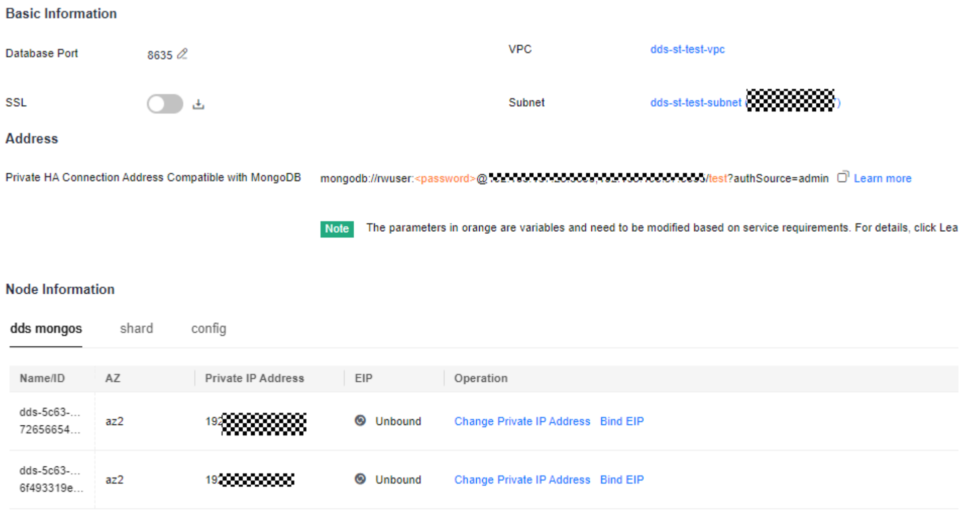

  Basic Information

  Database Port 8635

  VPC dds-st-test-vpc

  SSL

  Subnet dds-st-test-subnet ([redacted])

  Address

  Private HA Connection Address Compatible with MongoDB mongodb://rwuser:<password>@[redacted]/test?authSource=admin Learn more

  Note The parameters in orange are variables and need to be modified based on service requirements. For details, click Lea

  Node Information

Name/ID	AZ	Private IP Address	EIP	Operation
dds-5c63-... 72656654...	az2	19[redacted]	Unbound	Change Private IP Address Bind EIP
dds-5c63-... 6f493319e...	az2	19[redacted]	Unbound	Change Private IP Address Bind EIP

- **DB_PORT** is the port of the instance to be connected. The default port is 8635.

  Click the instance name. On the **Basic Information** page, choose **Connections**. On the **Private Connection** tab, obtain the database port information in the **Database Port** field in the **Basic Information** area.

  Figure 2-13 Obtaining the port

  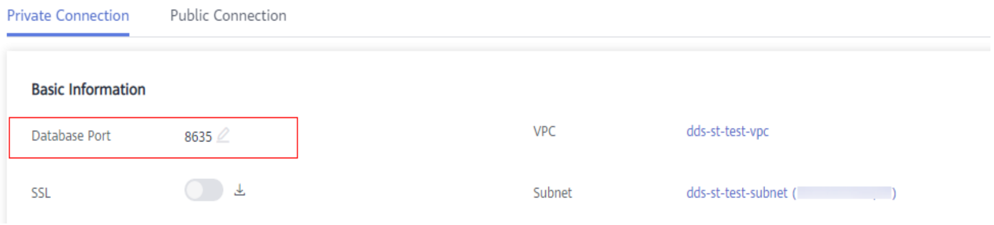

  Basic Information

  Database Port 8635

  VPC dds-st-test-vpc

  SSL

  Subnet dds-st-test-subnet ([redacted])

- **DB_USER** is the database user. The default value is **rwuser**.

Enter the database password when prompted:

```
Enter password:
```

Command example:

./mongo --host 192.168.1.6 --port 8635 -u rwuser -p --authenticationDatabase admin

Step 3 Check the connection result. If the following information is displayed, the connection is successful.

```
mongos>
```

----End

2.1.3.3 Connecting to Read Replicas Using Mongo Shell

Mongo shell is the default client for the MongoDB database server. You can use Mongo Shell to connect to DB instances, and query, update, and manage data in databases. DDS is compatible with MongoDB. Mongo Shell is a part of the MongoDB client. To use Mongo Shell, download and install the MongoDB client first, and then use the Mongo shell to connect to the DB instance.

By default, a DDS instance provides a private IP address. If your applications are deployed on an ECS and are in the same region and VPC as DDS instances, you can connect to DDS instances using a private IP address to achieve a fast transmission rate and high security.

This section describes how to use Mongo Shell to connect to a read replica over a private network.

You can connect to a read replica using an SSL connection or an unencrypted connection. The SSL connection is encrypted and more secure. To improve data transmission security, connect to instances using SSL.

Prerequisites

1. For details about how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS**.
2. Install the MongoDB client on the ECS. To ensure successful authentication, install the MongoDB client of the same version as the target instance.

   For details about how to install a MongoDB client, see **How Can I Install a MongoDB Client?**
3. The ECS can communicate with the DDS instance. For details, see **Configuring Security Group Rules**.

SSL Connection

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

Step 1 On the **Instances** page, click the instance name.

Step 2 In the navigation pane on the left, choose **Connections**.

Step 3 In the **Basic Information** area, click ⤓ next to the **SSL** field.

Step 4 Upload the root certificate to the ECS to be connected to the instance.

The following describes how to upload the certificate to a Linux and Window ECS:

- In Linux, run the following command:

  scp *<IDENTITY_FILE>* *<REMOTE_USER>***@***<REMOTE_ADDRESS>***:***<REMOTE_DIR>*

  NOTE

  – **IDENTITY_FILE** is the directory where the root certificate resides. The file access permission is 600.
  – **REMOTE_USER** is the ECS OS user.
  – **REMOTE_ADDRESS** is the ECS address.
  – **REMOTE_DIR** is the directory of the ECS to which the root certificate is uploaded.

- In Windows, upload the root certificate using a remote connection tool.

Step 5 Connect to a DDS instance. The DDS console provides the read replica connection address. You can use this address to connect to the read replica.

Example command:

./mongo "*<Read replica connection address>***" --ssl --sslCAFile** *<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **Read Replica Connection Address**: On the **Instances** page, click the instance to go to the **Basic Information** page. Choose **Connections**. Click the **Private Connection** tab. In the **Address** area, obtain the connection address of the read replica instance.

Figure 2-14 Obtaining the read replica connection address

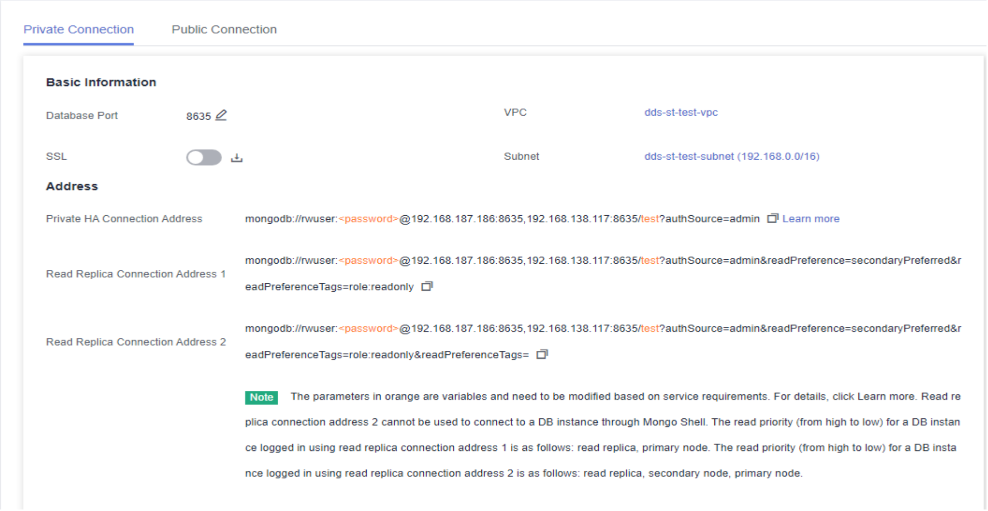

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet (192.168.0.0/16)

Address

Private HA Connection Address mongodb://rwuser:<password>@192.168.187.186:8635,192.168.138.117:8635/test?authSource=admin Learn more

Read Replica Connection Address 1 mongodb://rwuser:<password>@192.168.187.186:8635,192.168.138.117:8635/test?authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role:readonly

Read Replica Connection Address 2 mongodb://rwuser:<password>@192.168.187.186:8635,192.168.138.117:8635/test?authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role:readonly&readPreferenceTags=

Note The parameters in orange are variables and need to be modified based on service requirements. For details, click Learn more. Read replica connection address 2 cannot be used to connect to a DB instance through Mongo Shell. The read priority (from high to low) for a DB instance logged in using read replica connection address 1 is as follows: read replica, primary node. The read priority (from high to low) for a DB instance logged in using read replica connection address 2 is as follows: read replica, secondary node, primary node.

The format of the read replica connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:<password>@*192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role:readonly**

Pay attention to the following parameters in the read replica connection address:

Table 2-7 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.xx.xx:8635,192.168.xx.xx:8635	IP address and port of the mongos node of the cluster instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: To ensure that the internal communication of the cluster does not occupy resources such as the user IP address and bandwidth, the cluster certificate is generated using the internal management IP address. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

Command example:

./mongo "mongodb://rwuser:<password>@*192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role:readonly" --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

NOTE

When connecting to an instance using the read replica connection address, add double quotation marks (") before and after the connection information.

If the following information is displayed, the instance is successfully connected:

```
mongos>
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Log in to the ECS.

Step 2 Connect to a DDS instance. The DDS console provides the read replica connection address. You can use this address to connect to the read replica.

Example command:

./mongo "*<Read replica connection address>***"**

Read Replica Connection Address: On the **Instances** page, click the instance to go to the **Basic Information** page. Choose **Connections**. Click the **Private Connection** tab. In the **Address** area, obtain the connection address of the read replica instance.

Figure 2-15 Obtaining the read replica connection address

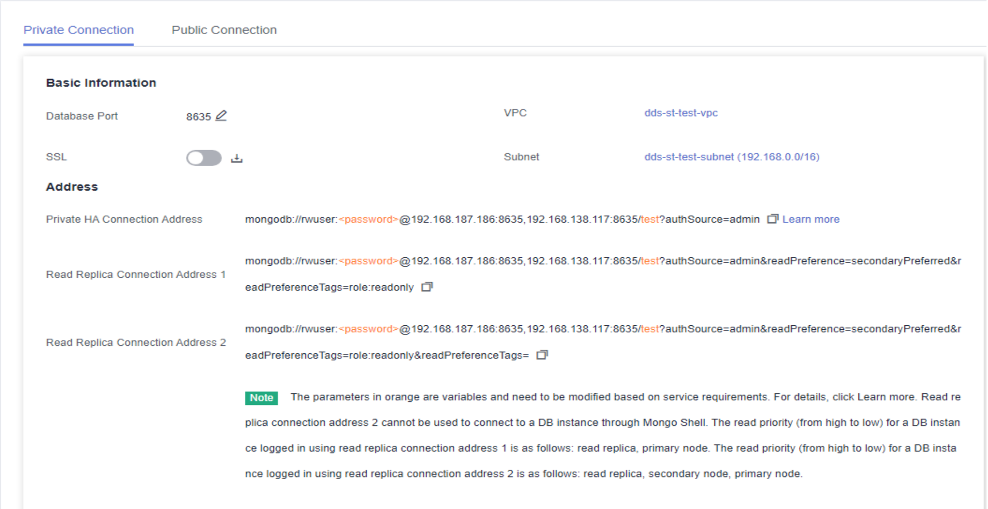

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet (192.168.0.0/16)

Address

Private HA Connection Address mongodb://rwuser:<password>@192.168.187.186:8635,192.168.138.117:8635/test?authSource=admin

Read Replica Connection Address 1 mongodb://rwuser:<password>@192.168.187.186:8635,192.168.138.117:8635/test?authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role:readonly

Read Replica Connection Address 2 mongodb://rwuser:<password>@192.168.187.186:8635,192.168.138.117:8635/test?authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role:readonly&readPreferenceTags=

Note The parameters in orange are variables and need to be modified based on service requirements. For details, click Learn more. Read replica connection address 2 cannot be used to connect to a DB instance through Mongo Shell. The read priority (from high to low) for a DB instance logged in using read replica connection address 1 is as follows: read replica, primary node. The read priority (from high to low) for a DB instance logged in using read replica connection address 2 is as follows: read replica, secondary node, primary node.

The format of the read replica connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:<password>@*192.168.xx.xx:8635,192.168.xx.xx:8635*/**test?authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role:readonly**

Pay attention to the following parameters in the private HA address:

Table 2-8 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.xx.xx:8635,192.168.xx.xx:8635	IP address and port of the mongos node of the cluster instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

Command example:

./mongo "mongodb://rwuser:<password>@*192.168.xx.xx:8635,192.168.xx.xx:8635*/**test?authSource=admin&readPreference=secondaryPreferred&readPreferenceTags=role:readonly"**

If the following information is displayed, the instance is successfully connected:

```
mongos>
```

----End

2.1.4 Connecting to a Cluster Instance over a Public Network

2.1.4.1 Binding and Unbinding an EIP

After you create a Cluster instance, you can bind an EIP to it to allow external access. If later you want to prohibit external access, you can also unbind the EIP from the instance.

Precautions

- Deleting a bound EIP does not mean that the EIP is unbound.
- Before accessing a database, apply for an EIP on the VPC console. Then, add an inbound rule to allow the IP addresses or IP address ranges of ECSs. For details, see **Configuring a Security Group**.

- In the cluster instance, only dds mongos can have an EIP bound. To change the EIP that has been bound to a node, you need to unbind it from the node first.

Binding an EIP

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the navigation pane on the left, choose **Connections**. Click the **Public Connection** tab. In the **Basic Information** area, locate the dds mongos node and click **Bind EIP** in the **Operation** column.

Figure 2-16 Binding an EIP

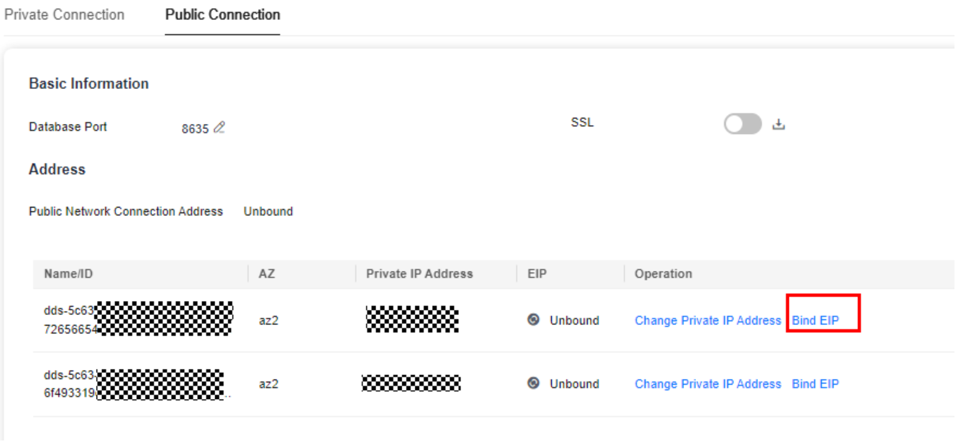

Private Connection | Public Connection

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address Unbound

Name/ID	AZ	Private IP Address	EIP	Operation
dds-5c63[redacted] 72656654[redacted]	az2	[redacted]	Unbound	Change Private IP Address Bind EIP
dds-5c63[redacted] 6f493319[redacted]	az2	[redacted]	Unbound	Change Private IP Address Bind EIP

Alternatively, in the **Node Information** area on the **Basic Information** page, locate the dds mongos node and choose **More** > **Bind EIP** in the **Operation** column.

Figure 2-17 Binding an EIP

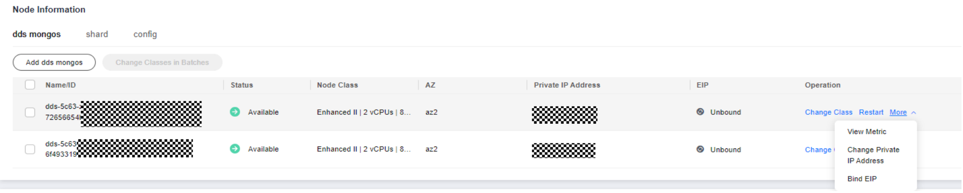

Node Information

dds mongos | shard | config

Add dds mongos | Change Classes in Batches

□	Name/ID	Status	Node Class	AZ	Private IP Address	EIP	Operation
□	dds-5c63-[redacted] 72656654[redacted]	Available	Enhanced II | 2 vCPUs | 8...	az2	[redacted]	Unbound	Change Class Restart More
□	dds-5c63[redacted] 6f493319[redacted]	Available	Enhanced II | 2 vCPUs | 8...	az2	[redacted]	Unbound	Change

Step 6 In the displayed dialog box, all available unbound EIPs are listed. Select the required EIP and click **OK**. If no available EIPs are displayed, click **View EIP** and create an EIP on the VPC console.

Figure 2-18 Selecting an EIP

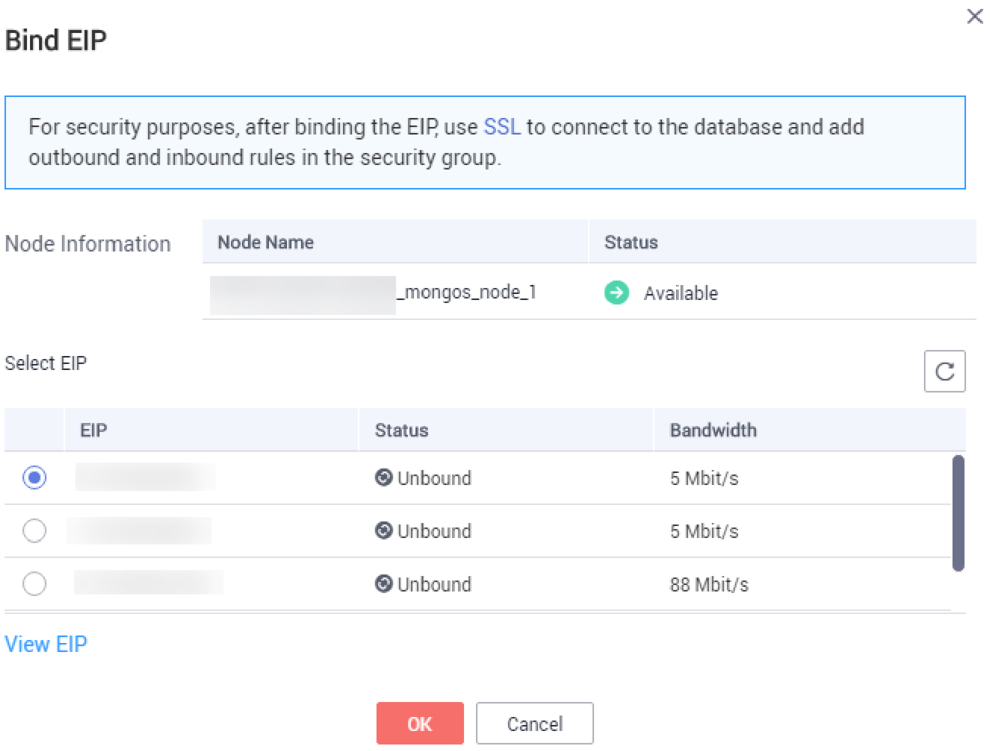

Bind EIP

For security purposes, after binding the EIP, use SSL to connect to the database and add outbound and inbound rules in the security group.

Node Information

Node Name	Status
[redacted]_mongos_node_1	Available

Select EIP

	EIP	Status	Bandwidth
●	[redacted]	Unbound	5 Mbit/s
○	[redacted]	Unbound	5 Mbit/s
○	[redacted]	Unbound	88 Mbit/s

Step 7 In the **EIP** column on the **dds mongos** tab, you can view the EIP that was bound.

To unbind an EIP from the instance, see **Unbinding an EIP**.

----End

Unbinding an EIP

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the navigation pane on the left, choose **Connections**. Click the **Public Connection** tab. In the **Basic Information** area, locate the dds mongos node and click **Unbind EIP** in the **Operation** column.

Figure 2-19 Unbinding an EIP

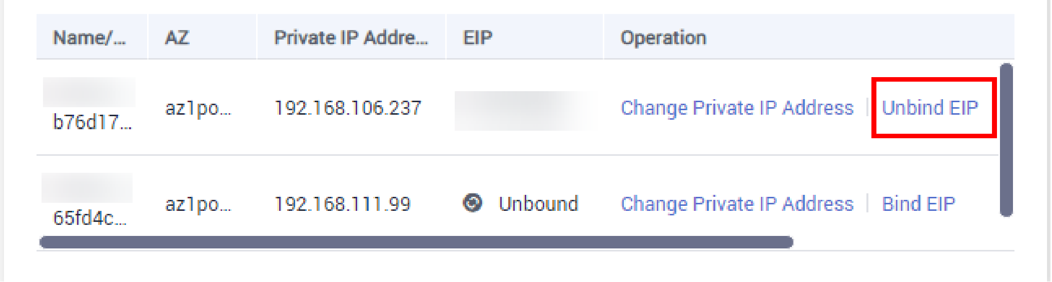

Name/...	AZ	Private IP Addre...	EIP	Operation
[redacted] b76d17...	az1po...	192.168.106.237	[redacted]	Change Private IP Address | Unbind EIP
[redacted] 65fd4c...	az1po...	192.168.111.99	Unbound	Change Private IP Address | Bind EIP

Alternatively, in the **Node Information** area on the **Basic Information** page, locate the dds mongos node and choose **More** > **Unbind EIP** in the **Operation** column.

Figure 2-20 Unbinding an EIP

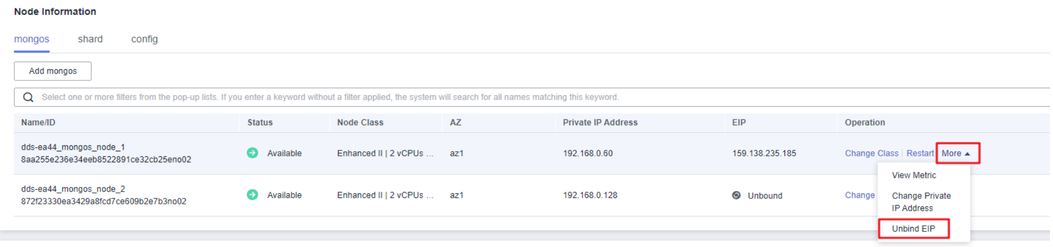

Node Information

mongos
shard
config
Add mongos
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Status	Node Class	AZ	Private IP Address	EIP	Operation
dds-ea44_mongos_node_1 8aa255e236e34eeb8522891ce32cb25eno02	Available	Enhanced II | 2 vCPUs ...	az1	192.168.0.60	159.138.235.185	Change Class | Restart | More ▴
dds-ea44_mongos_node_2 872f23330ea3429a8fcd7ce609b2e7b3no02	Available	Enhanced II | 2 vCPUs ...	az1	192.168.0.128	Unbound	Change

View Metric
Change Private IP Address
Unbind EIP

Step 6 In the displayed dialog box, click **Yes**.

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

----End

2.1.4.2 Configuring a Security Group

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

To ensure database security and reliability, you need to configure security group rules to allow specific IP addresses and ports to access DDS instances.

To access an instance from the Internet, add an inbound rule for the security group associated with the instance.

Precautions

- By default, an account can create up to 500 security group rules.
- Too many security group rules will increase the first packet latency, so a maximum of 50 rules for each security group is recommended.
- By default, one DDS instance is associated with only one security group.
- DDS allows you to associate multiple security groups to a DB instance. You can apply for the service based on your service requirements. For better network performance, you are advised to select no more than five security groups.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Network Information** area on the **Basic Information** page, click the security group.

Figure 2-21 Security Group

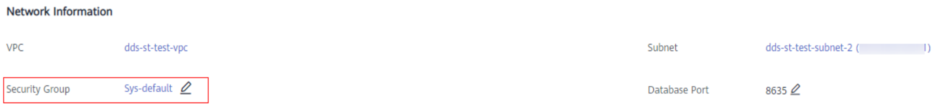

Network Information

VPC dds-st-test-vpc

Subnet dds-st-test-subnet-2 ([redacted])

Security Group Sys-default

Database Port 8635

You can also choose **Connections** in the navigation pane on the left. On the **Public Connection** tab, in the **Security Group** area, click the security group name.

Figure 2-22 Security Group

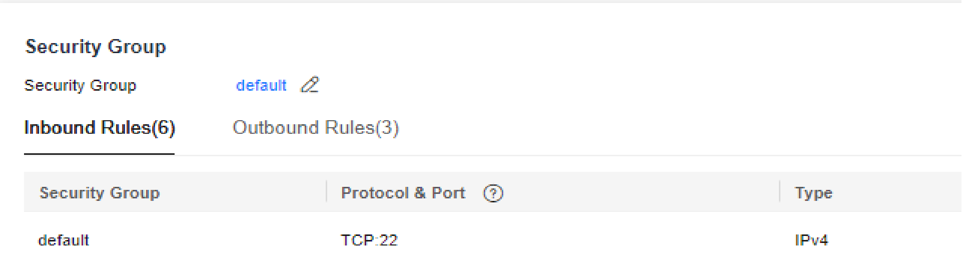

Security Group

Security Group default

Security Group	Protocol & Port	Type
default	TCP:22	IPv4

Step 6 On the **Security Group** page, locate the target security group and click **Manage Rule** in the **Operation** column.

Step 7 On the **Inbound Rules** tab, click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Step 8 Add a security group rule as prompted.

Figure 2-23 Add Inbound Rule

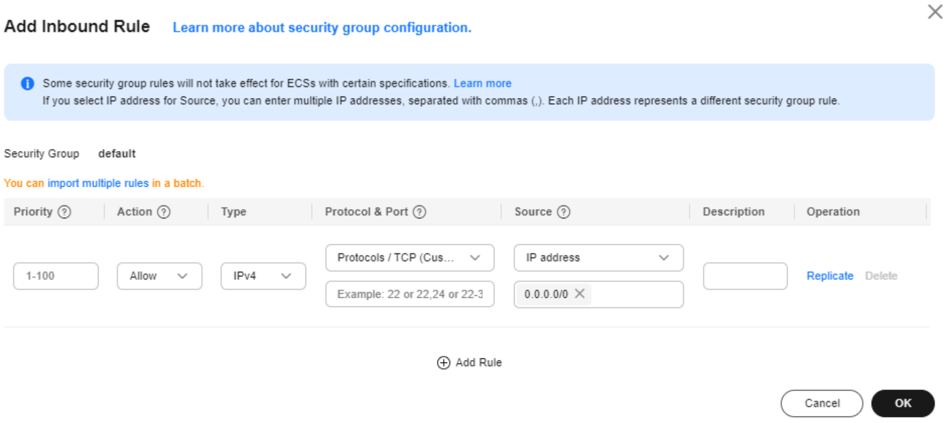

Add Inbound Rule Learn more about security group configuration.

Some security group rules will not take effect for ECSs with certain specifications. Learn more
If you select IP address for Source, you can enter multiple IP addresses, separated with commas (,). Each IP address represents a different security group rule.

Security Group default

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Source	Description	Operation
1-100	Allow	IPv4	Protocols / TCP (Cus... Example: 22 or 22,24 or 22-3	IP address 0.0.0.0/0		Replicate Delete

Table 2-9 Inbound rule settings

Parameter	Description	Example
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default priority is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	The security group rule actions. A rule with a deny action overrides another with an allow action if the two rules have the same priority.	Allow
Protocol & Port	The network protocol required for access. The option can be **All**, **TCP**, **UDP**, **ICMP**, or **GRE**.	TCP
	Port: the port on which you wish to allow access to DDS. The default port is 8635. The port ranges from 2100 to 9500 or can be 27017, 27018, or 27019.	8635
Type	IP address type. Only **IPv4** and **IPv6** are supported.	IPv4

Parameter	Description	Example
Source	Specifies the supported IP address, security group, and IP address group, which allow access from IP addresses or instances in other security group. Example: ● Single IP address: 192.168.10.10/32 ● IP address segment: 192.168.1.0/24 ● All IP addresses: 0.0.0.0/0 ● Security group: sg-abc ● IP address group: ipGroup-test If you enter a security group, all ECSs associated with the security group comply with the created rule. For more information about IP address groups, see **IP Address Group Overview**.	0.0.0.0/0
Description	(Optional) Provides supplementary information about the security group rule. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

Step 9 Click **OK**.

----End

2.1.4.3 Connecting to a Cluster Instance Using Mongo Shell (Public Network)

In the following scenarios, you can access a DDS instance from the Internet by binding an EIP to the instance.

Scenario 1: Your applications are running on an ECS that is in a different region from the one where the DDS instance is located.

Figure 2-24 Accessing DDS from ECS across regions

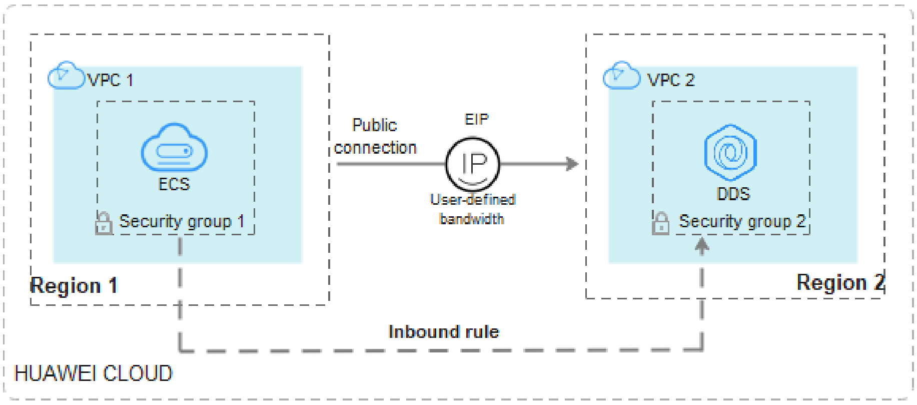

Scenario 2: Your applications are deployed on a cloud server provided by other vendors.

Figure 2-25 Accessing DDS from other cloud servers

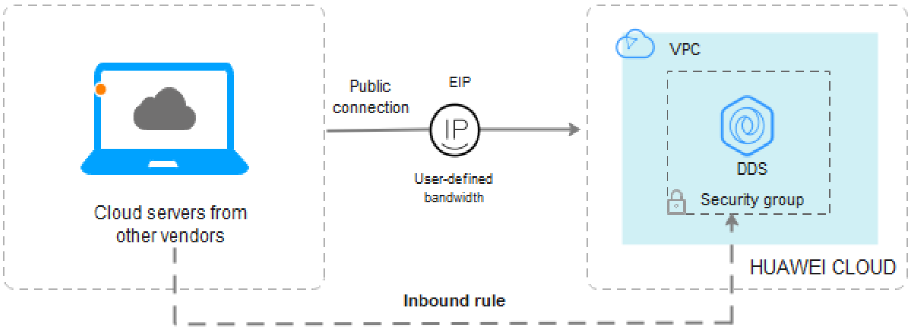

This section describes how to use Mongo Shell to connect to a cluster instance over a public network.

You can connect to an instance using an SSL connection or an unencrypted connection. The SSL connection is encrypted and more secure. To improve data transmission security, connect to instances using SSL.

Prerequisites

1. For details about how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS**.
2. **Bind an EIP** to the cluster instance and **set security group rules** to ensure that the instance can be accessed from the ECS.
3. You have installed the MongoDB client on the ECS.

   For details about how to install a MongoDB client, see **How Can I Install a MongoDB Client?**

SSL

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

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

Step 2 Click ⊙ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Connections**.

Step 6 In the **Basic Information** area, click ⤓ next to the **SSL** field.

Step 7 Upload the root certificate obtained in **Step 6** to the ECS.

The following describes how to upload the certificate to a Linux and Windows ECS:

- In Linux, run the following command:

  scp*<IDENTITY_FILE><REMOTE_USER>***@***<REMOTE_ADDRESS>***:***<REMOTE_DIR>*

  NOTE

  - **IDENTITY_FILE** is the directory where the root certificate resides. The file access permission is 600.
  - **REMOTE_USER** is the ECS OS user.
  - **REMOTE_ADDRESS** is the ECS address.
  - **REMOTE_DIR** is the directory of the ECS to which the root certificate is uploaded.

- In Windows, upload the root certificate using a remote connection tool.

Step 8 Connect to the instance in the directory where the MongoDB client is located.

Method 1: Using a public network connection address

Example command:

./mongo *<Public network connection address>* **--ssl --sslCAFile** *<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **Public Network Connection Address**: On the **Instances** page, click the instance to switch to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Public Connection** tab. In the **Address** area, obtain the instance connection address from the **Public Network Connection Address** field.

Figure 2-26 Obtaining the public network connection address

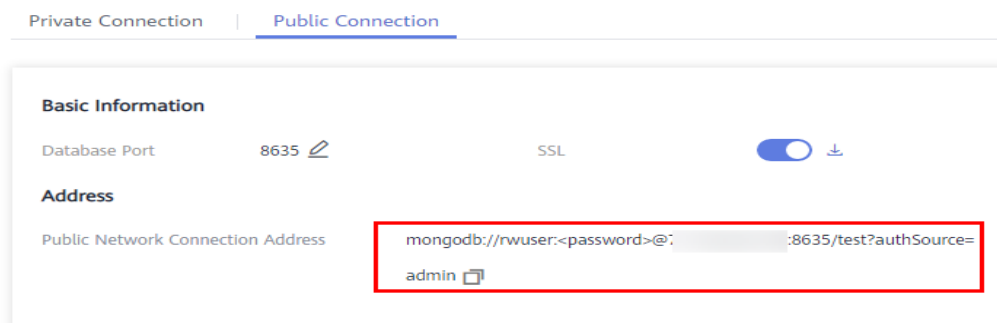

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address

mongodb://rwuser:<password>@[redacted]:8635/test?authSource=admin

The format of the public connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:*<password>*@192.168.*xx.xx*:8635**/test?authSource=admin**

Pay attention to the following parameters in the public connection address:

Table 2-10 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.*xx.xx*:8635	EIP and port bound to the dds mongos node of the cluster instance
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: To ensure that the internal communication of the cluster does not occupy resources such as the user IP address and bandwidth, the cluster certificate is generated using the internal management IP address. **--sslAllowInvalidHostnames** is needed for the SSL connection through a public network.

Command example:

./mongo mongodb://rwuser:*<password>*@192.168.*xx.xx*:8635**/test?authSource=admin --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

Method 2: Using an EIP

Example command:

./mongo --host <*DB_HOST*> **--port** <*DB_PORT*> **-u** <*DB_USER*> **-p --authenticationDatabase admin --ssl --sslCAFile** <*FILE_PATH*> **--sslAllowInvalidHostnames**

Parameter description:

- **DB_HOST** is the EIP bound to the instance to be connected.

  You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the **Public Connection** tab, obtain the EIP bound to the dds mongos node in the **EIP** column.

  If there are multiple dds mongos nodes, the EIP of any node can be used to connect to the instance.

  Figure 2-27 Obtaining an EIP

  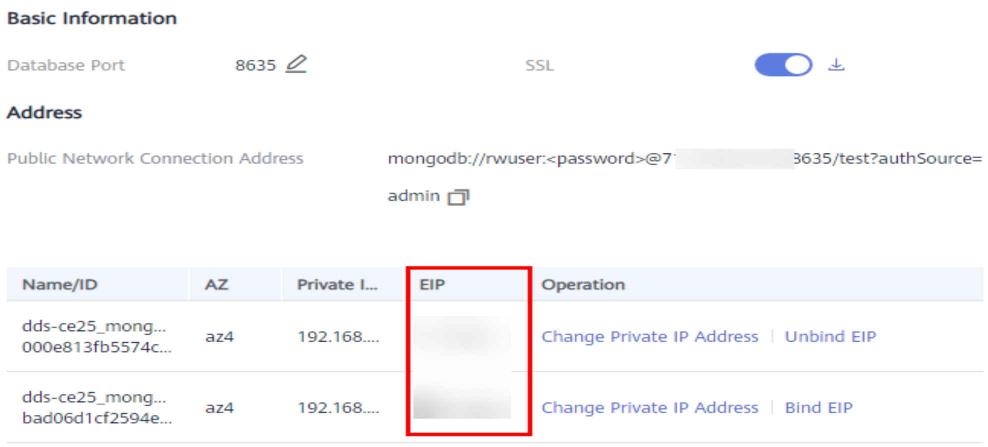

  Basic Information

  Database Port 8635 SSL

  Address

  Public Network Connection Address mongodb://rwuser:<password>@7[redacted]3635/test?authSource=admin

Name/ID	AZ	Private I...	EIP	Operation
dds-ce25_mong... 000e813fb5574c...	az4	192.168....	[redacted]	Change Private IP Address | Unbind EIP
dds-ce25_mong... bad06d1cf2594e...	az4	192.168....	[redacted]	Change Private IP Address | Bind EIP

- **DB_PORT** is the port of the instance to be connected. The default port number is 8635.

  You can click the instance to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Public Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

  Figure 2-28 Obtaining the port

  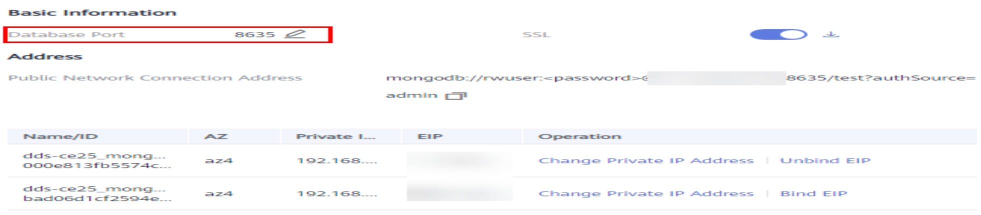

  Basic Information

  Database Port 8635 SSL

  Address

  Public Network Connection Address mongodb://rwuser:<password>@[redacted]8635/test?authSource=admin

Name/ID	AZ	Private I...	EIP	Operation
dds-ce25_mong... 000e813fb5574c...	az4	192.168....	[redacted]	Change Private IP Address | Unbind EIP
dds-ce25_mong... bad06d1cf2594e...	az4	192.168....	[redacted]	Change Private IP Address | Bind EIP

- **DB_USER** is the database user. The default value is **rwuser**.
- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: To ensure that the internal communication of the cluster does not occupy resources such as the user IP address and

bandwidth, the cluster certificate is generated using the internal management IP address. **--sslAllowInvalidHostnames** is needed for the SSL connection through a public network.

Enter the database account password when prompted:

```
Enter password:
```

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

Step 9 Check the connection result. If the following information is displayed, the connection is successful.

```
mongos>
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Log in to the ECS.

Step 2 Connect to the instance in the directory where the MongoDB client is located.

Method 1: Using a public network connection address

Example command:

./mongo *<Public network address>*

Public Network Connection Address: You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Public Connection** tab. In the **Address** area, obtain the instance connection address from the **Public Network Connection Address** field.

Figure 2-29 Obtaining the public network connection address

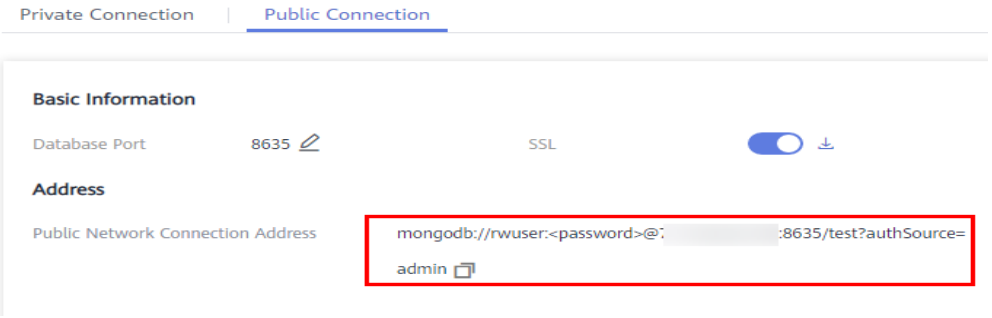

Private Connection | Public Connection

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address mongodb://rwuser:<password>@[redacted]:8635/test?authSource=admin

The format of the public connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:*<password>*@192.168.*xx.xx*:8635**/test?authSource=admin**

The following table describes the required parameters in the public connection address.

Table 2-11 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.*xx.xx*:8635	EIP and port bound to the dds mongos node of the cluster instance
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

Command example:

./mongo mongodb://rwuser:*<password>*@192.168.*xx.xx*:8635**/test?authSource=admin**

Method 2: Using an EIP

Example command:

./mongo --host <*DB_HOST*> **--port** <*DB_PORT*> **-u** <*DB_USER*> **-p --authenticationDatabase admin**

Parameter description:

- **DB_HOST** is the EIP bound to the instance to be connected.

  You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the **Public Connection** tab, obtain the EIP bound to the dds mongos node in the **EIP** column.

  If there are multiple dds mongos nodes, the EIP of any node can be used to connect to the instance.

Figure 2-30 Obtaining an EIP

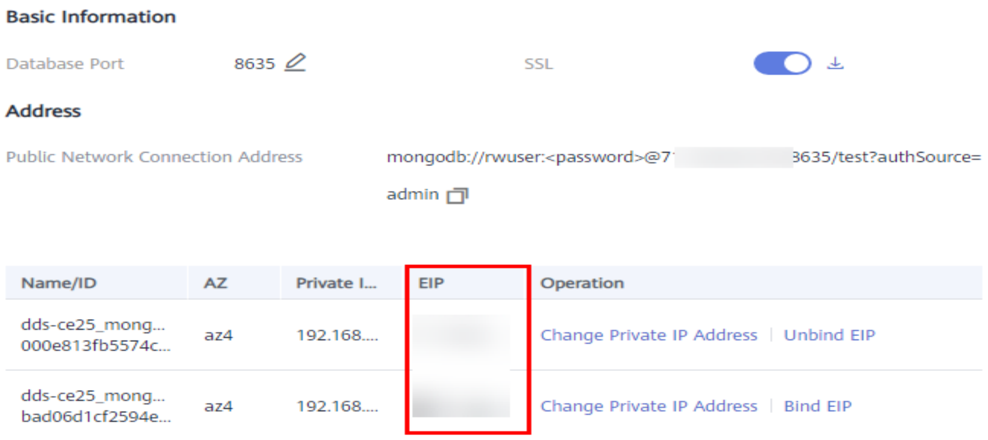

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address mongodb://rwuser:<password>@7[redacted]3635/test?authSource=admin

Name/ID	AZ	Private I...	EIP	Operation
dds-ce25_mong... 000e813fb5574c...	az4	192.168....	[redacted]	Change Private IP Address | Unbind EIP
dds-ce25_mong... bad06d1cf2594e...	az4	192.168....	[redacted]	Change Private IP Address | Bind EIP

- **DB_PORT** is the port of the instance to be connected. The default port number is 8635.

  You can click the instance to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Public Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

Figure 2-31 Obtaining the port

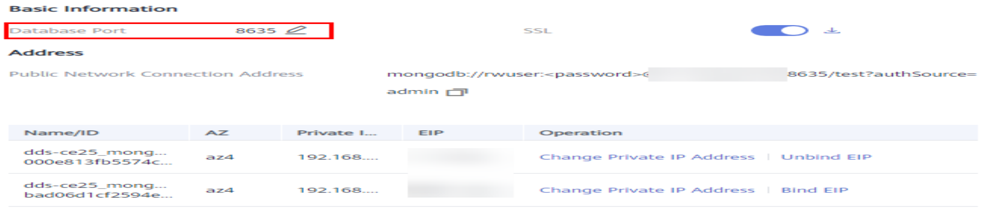

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address mongodb://rwuser:<password>@[redacted]8635/test?authSource=admin

Name/ID	AZ	Private I...	EIP	Operation
dds-ce25_mong... 000e813fb5574c...	az4	192.168....	[redacted]	Change Private IP Address | Unbind EIP
dds-ce25_mong... bad06d1cf2594e...	az4	192.168....	[redacted]	Change Private IP Address | Bind EIP

- **DB_USER** is the database user. The default value is **rwuser**.

Enter the database account password when prompted:

```
Enter password:
```

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin**

Step 3 Check the connection result. If the following information is displayed, the connection is successful.

```
mongos>
```

----End

2.1.4.4 Connecting to a Cluster Instance Using Robo 3T

To connect to an instance from a local device, you can use Robo 3T to access the instance from the Internet.

This section describes how to use Robo 3T to connect to a cluster instance from a local device. In this section, the Windows operating system (OS) used by the client is used as an example.

Robo 3T can connect to an instance with an unencrypted connection or an encrypted connection (SSL). To improve data transmission security, connect to instances using SSL.

Connection Diagram

Figure 2-32 Connection diagram

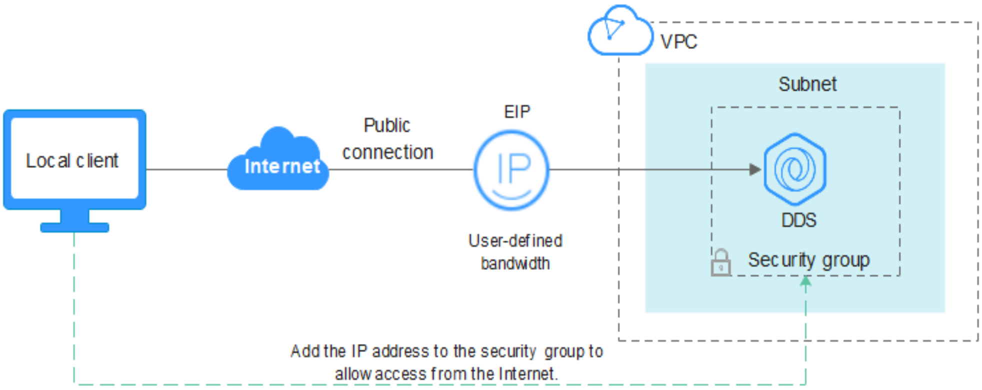

Prerequisites

1. **Bind an EIP** to the cluster instance and **configure security group rules** to ensure that the instance can be accessed using Robo 3T.
2. Install Robo 3T.

   For details, see **Installing Robo 3T**.

SSL

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

Step 1 Run the installed Robo 3T. On the displayed dialog box, click **Create**.

Figure 2-33 Connections

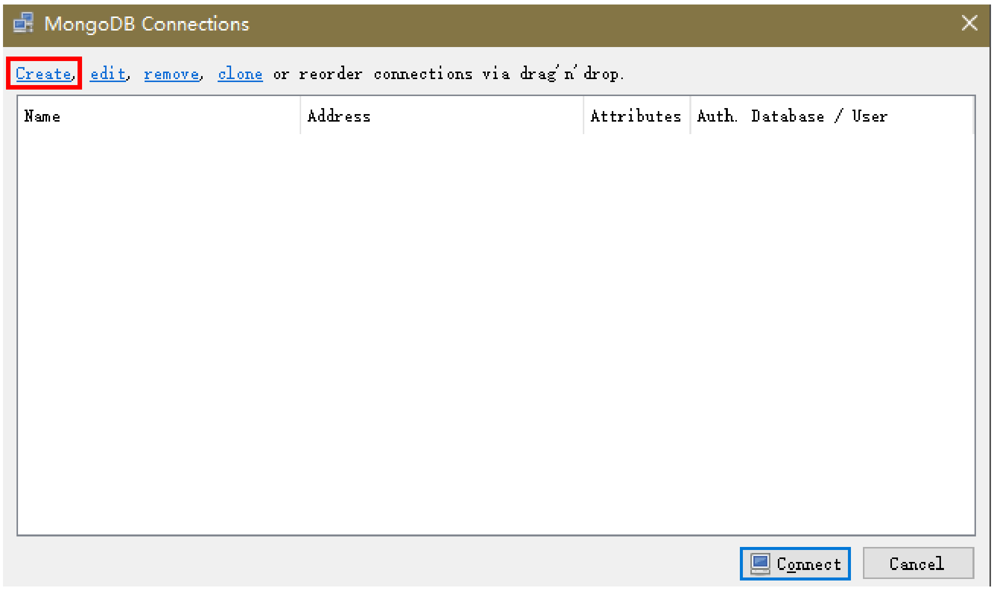

Step 2 In the **Connection Settings** dialog box, set the parameters of the new connection.

1. On the **Connection** tab, enter the name of the new connection in the **Name** text box and enter the EIP and database port that are bound to the DDS DB instance in the **Address** text box.

Figure 2-34 Connection

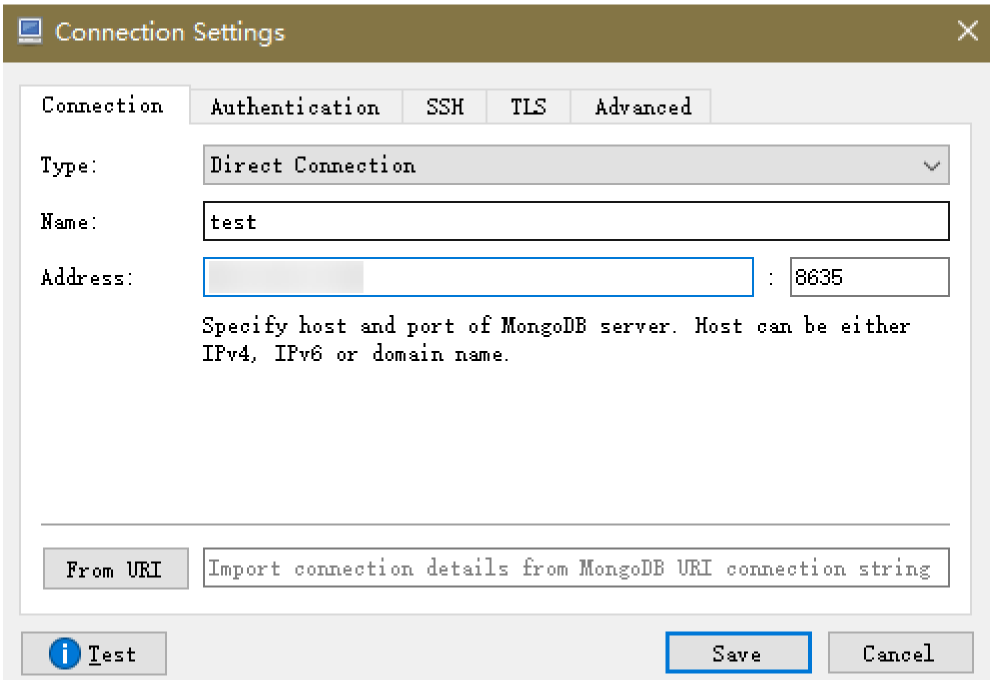

Type: Direct Connection

Name: test

Address: [redacted] : 8635

Specify host and port of MongoDB server. Host can be either IPv4, IPv6 or domain name.

2. On the **Authentication** tab, set **Database** to **admin**, **User Name** to **rwuser**, and **Password** to the administrator password you set during the creation of the cluster instance.

Figure 2-35 Authentication

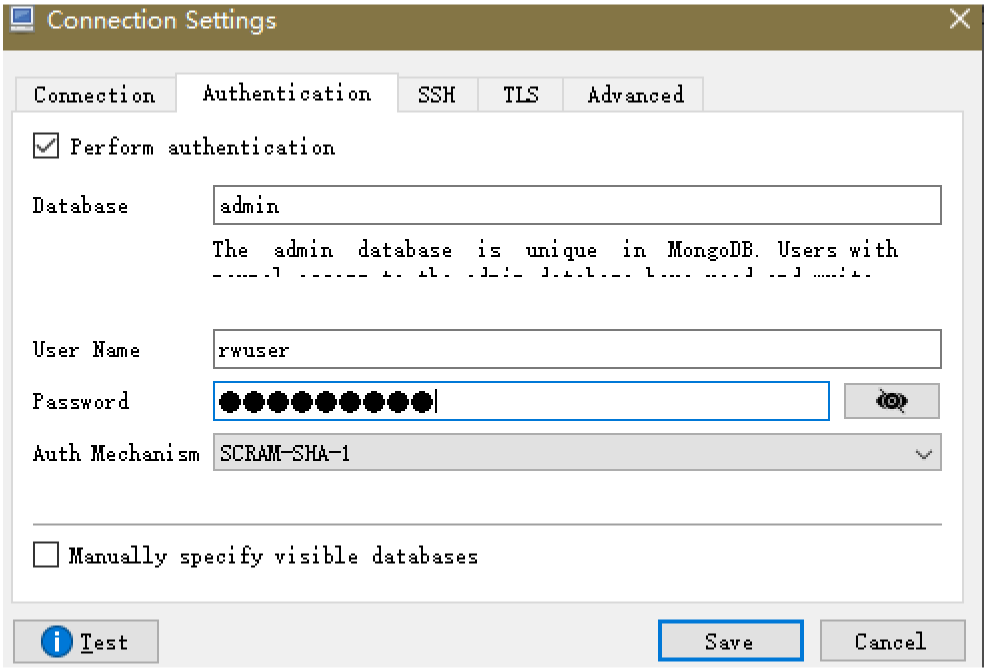

☑ Perform authentication

Database: admin

The admin database is unique in MongoDB. Users with

User Name: rwuser

Password: ●●●●●●●●●●

Auth Mechanism: SCRAM-SHA-1

□ Manually specify visible databases

3. On the **TLS** tab, select **Use TLS protocol** and select **Self-signed Certificate** for **Authentication Method**.

Figure 2-36 SSL

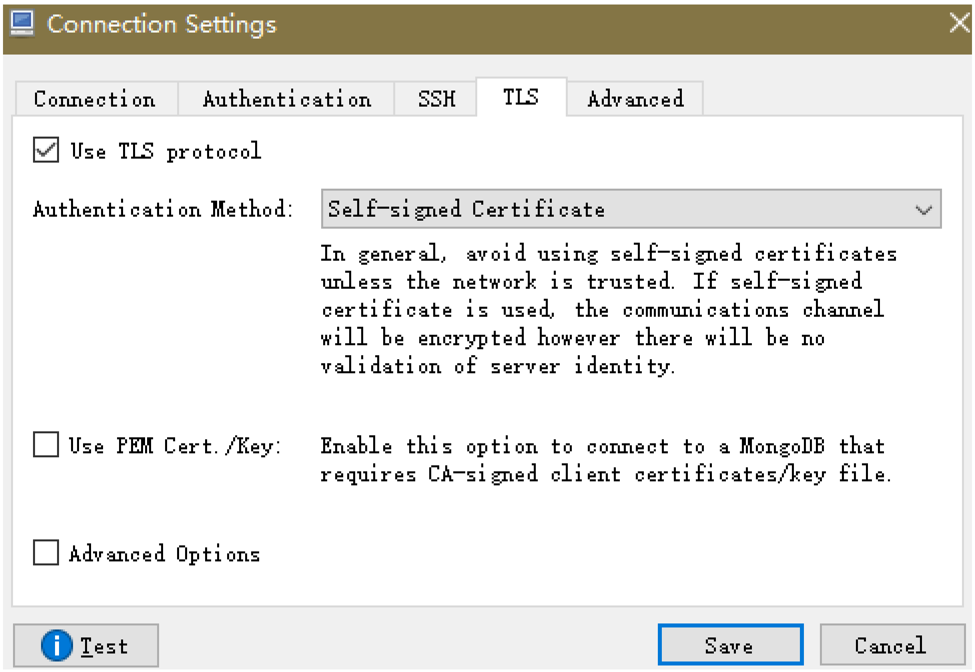

☑ Use TLS protocol

Authentication Method: Self-signed Certificate

In general, avoid using self-signed certificates unless the network is trusted. If self-signed certificate is used, the communications channel will be encrypted however there will be no validation of server identity.

□ Use PEM Cert./Key: Enable this option to connect to a MongoDB that requires CA-signed client certificates/key file.

□ Advanced Options

4. Click **Save**.

Step 3 On the **MongoDB Connections** page, click **Connect** to connect to the cluster instance.

Figure 2-37 Cluster connection information

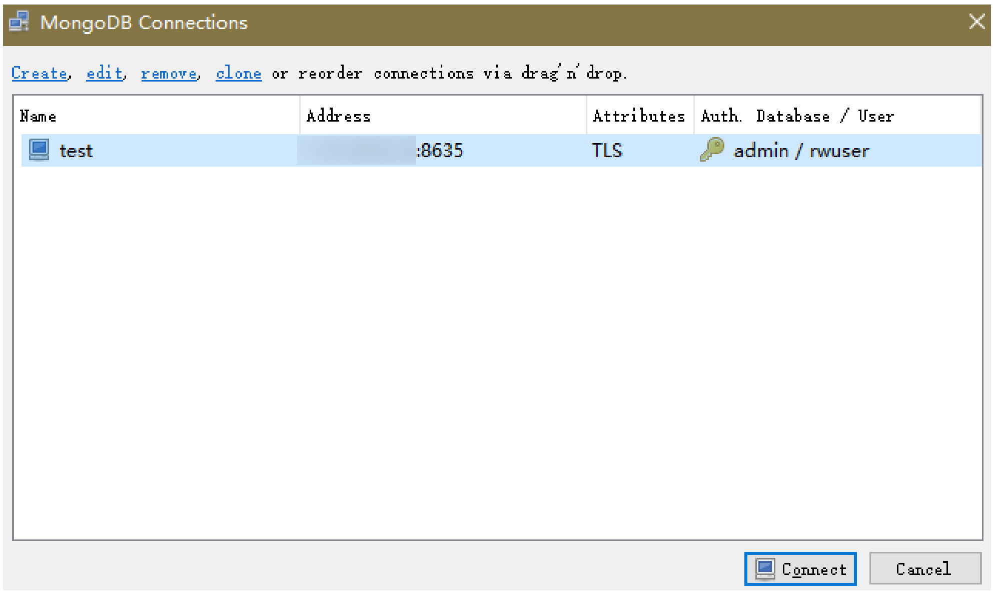
MongoDB Connections

Create, edit, remove, clone or reorder connections via drag'n'drop.

Name	Address	Attributes	Auth. Database / User
test	[redacted]:8635	TLS	admin / rwuser

Step 4 If the cluster instance is successfully connected, the page shown in **Figure 2-38** is displayed.

Figure 2-38 Cluster connected successfully

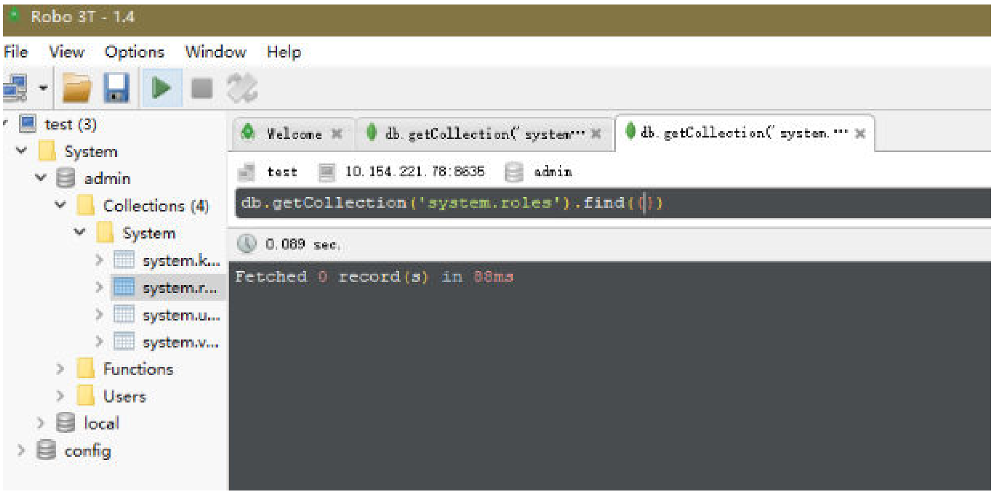

test   10.154.221.78:8635   admin

```
db.getCollection('system.roles').find({})
```

0.089 sec.

```
Fetched 0 record(s) in 88ms
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details, see **Enabling and Disabling SSL**.

Step 1 Run the installed Robo 3T. On the displayed dialog box, click **Create**.

Figure 2-39 Connections

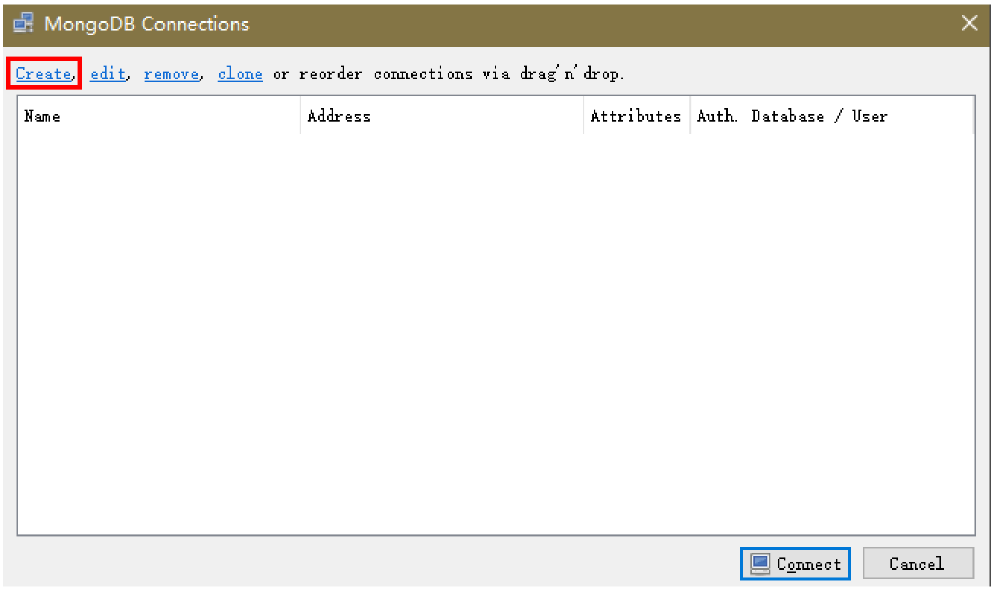

Create, edit, remove, clone or reorder connections via drag'n' drop.

Step 2 In the **Connection Settings** dialog box, set the parameters of the new connection.

1. On the **Connection** tab, enter the name of the new connection in the **Name** text box and enter the EIP and database port that are bound to the DDS DB instance in the **Address** text box.

Figure 2-40 Connection

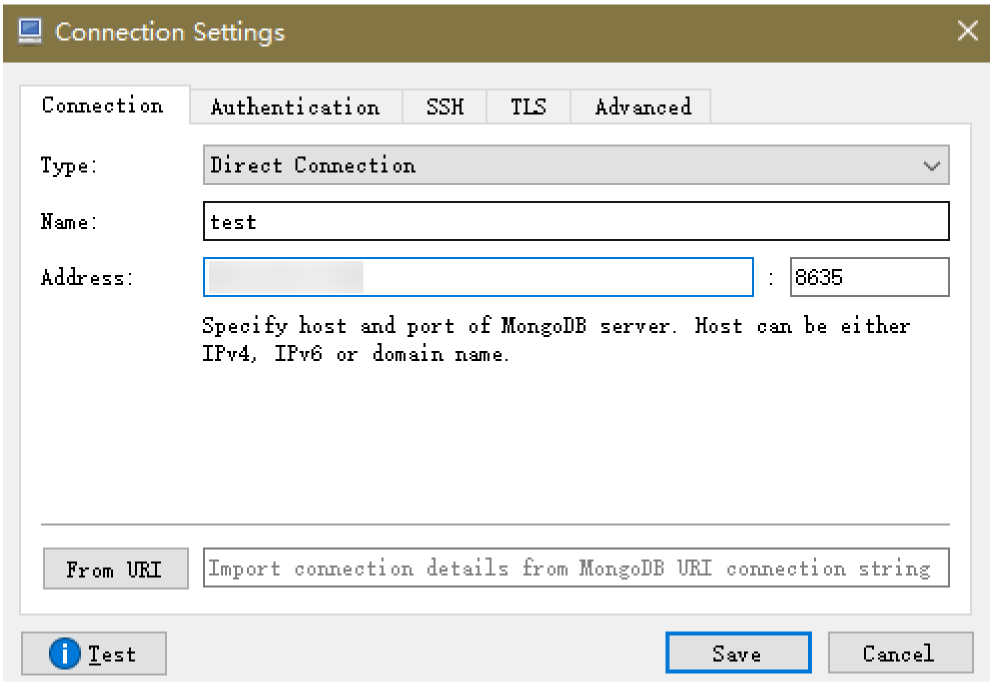

Type: Direct Connection

Name: test

Address: : 8635

Specify host and port of MongoDB server. Host can be either IPv4, IPv6 or domain name.

Import connection details from MongoDB URI connection string

2. On the **Authentication** tab, set **Database** to **admin**, **User Name** to **rwuser**, and **Password** to the administrator password you set during the creation of the cluster instance.

Figure 2-41 Authentication

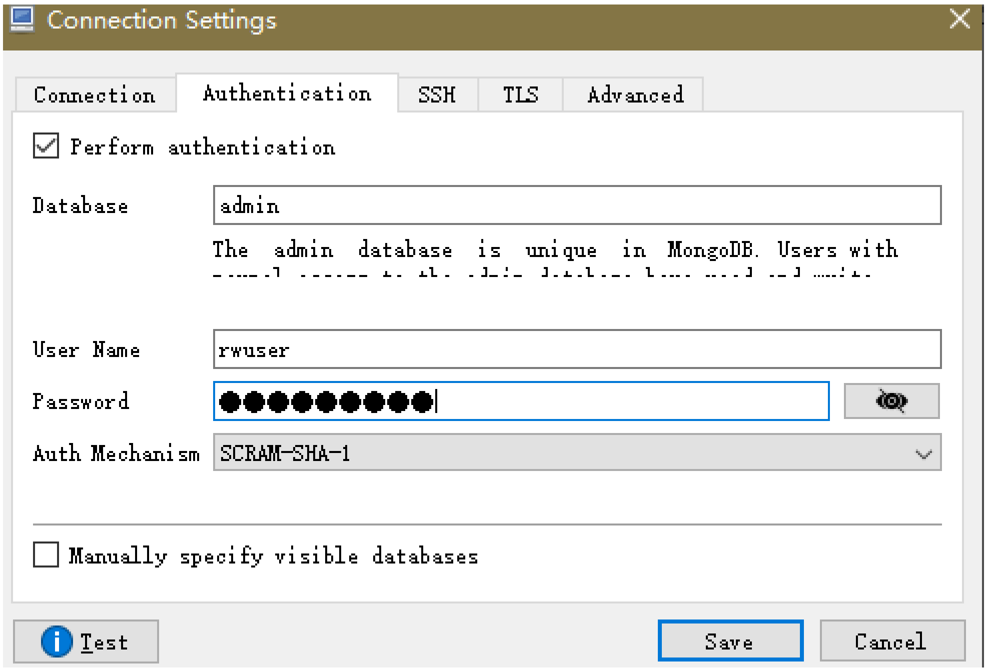

Connection Settings

☑ Perform authentication

Database: admin

The admin database is unique in MongoDB. Users with

User Name: rwuser

Password: ●●●●●●●●●●

Auth Mechanism: SCRAM-SHA-1

□ Manually specify visible databases

3. Click **Save**.

Step 3 On the **MongoDB Connections** page, click **Connect** to connect to the cluster instance.

Figure 2-42 Cluster connection information

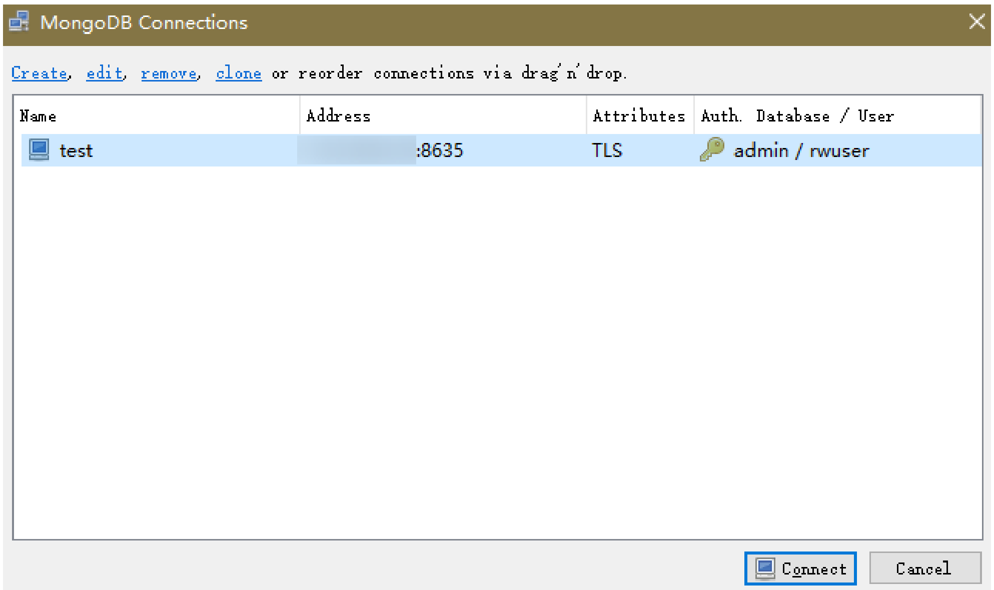

MongoDB Connections

Create, edit, remove, clone or reorder connections via drag'n'drop.

Name	Address	Attributes	Auth. Database / User
test	[redacted]:8635	TLS	admin / rwuser

Step 4 If the cluster instance is successfully connected, the page shown in **Figure 2-43** is displayed.

Figure 2-43 Cluster connected successfully

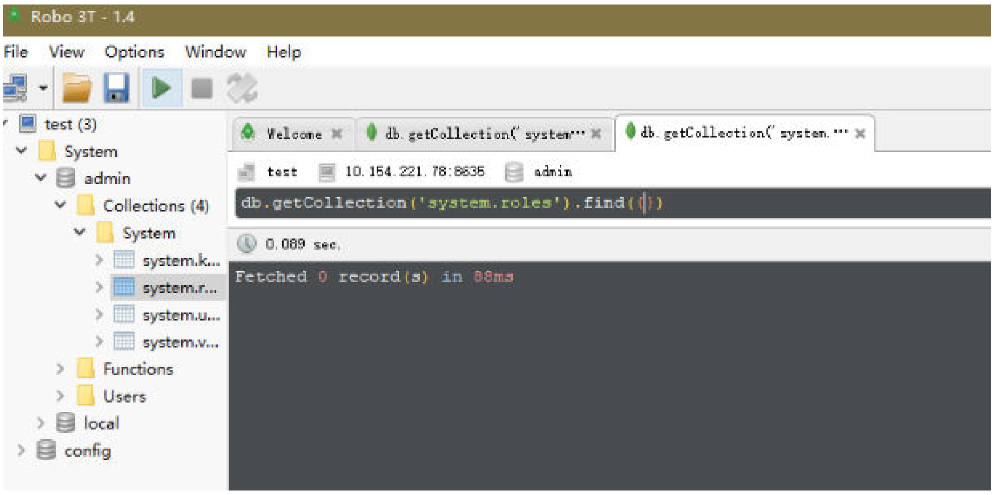

```
db.getCollection('system.roles').find({})
```

0.089 sec.

```
Fetched 0 record(s) in 88ms
```

----End

2.1.5 Connecting to a Cluster Instance Using Program Code

2.1.5.1 Java

If you are connecting to an instance using Java, an SSL certificate is optional, but downloading an SSL certificate and encrypting the connection will improve the security of your instance. SSL is disabled by default for newly created instances, but you can enable SSL by referring to **Enabling or Disabling SSL**. SSL encrypts connections to databases but it increases the connection response time and CPU usage. For this reason, enabling SSL is not recommended.

Prerequisites

Familiarize yourself with:

- Computer basics
- Java code

Obtaining and Using Java

- Download the Jar driver from: **https://repo1.maven.org/maven2/org/mongodb/mongo-java-driver/3.0.4/**
- To view the usage guide, visit **https://mongodb.github.io/mongo-java-driver/4.2/driver/getting-started/installation/**.

Using an SSL Certificate

NOTE

- Download the SSL certificate and verify the certificate before connecting to databases.
- On the **Instances** page, click the target DB instance name. In the **DB Information** area on the **Basic Information** page, click ⤓ in the **SSL** field to download the root certificate or certificate bundle.
- For details about how to set up an SSL connection, see the MongoDB Java Driver official document at **https://www.mongodb.com/docs/drivers/java/sync/current/fundamentals/connection/tls/#std-label-tls-ssl**.
- Java Runtime Environment (JRE) earlier than Java 8 enables TLS 1.2 only in updated versions. If TLS 1.2 is not enabled for your JRE, upgrade it to a later version to use TLS 1.2 for connection.

If you connect to a cluster instance using Java, the format of code is as follows:

```
mongodb://<username>:<password>@<instance_ip>:<instance_port>/<database_name>?authSource=admin&ssl=true
```

Table 2-12 Parameter description

Parameter	Description
<username>	Current username.
<password>	Password for the current username
<instance_ip>	If you attempt to access the instance from an ECS, set *instance_ip* to the private IP address displayed on the **Basic Information** page of the instance to which you intend to connect.
	If you intend to access the instance through an EIP, set *instance_ip* to the EIP that has been bound to the instance.
	If multiple host addresses are required, list the addresses in the format of <instance_ip1>:<instance_port1>,<instance_ip2>:<instance_port2>....... Example: mongodb://username:*****@127.***.***.1:8635,127.***.***.2:8635/?authSource=admin
<instance_port>	Database port displayed on the **Basic Information** page. Default value: **8635**
<database_name>	Name of the database to be connected.
authSource	Authentication user database. The value is **admin**.
ssl	Connection mode. **true** indicates that the SSL connection mode is used.

Use the keytool to configure the CA certificate. For details about the parameters, see **Table 2-13**.

```
keytool -importcert -trustcacerts -file <path to certificate authority file> -keystore <path to trust store> -
storepass <password>
```

Table 2-13 Parameter description

Parameter	Description
<path to certificate authority file>	Path for storing the SSL certificate.
<path to trust store>	Path for storing the truststore. Set this parameter as required, for example, **./trust/certs.keystore**.
<password>	Custom password.

Set the JVM system properties in the program to point to the correct truststore and keystore:

- System.setProperty("javax.net.ssl.trustStore","<path to trust store>");
- System.setProperty("javax.net.ssl.trustStorePassword","<password>");

For details about the Java code, see the following example:

```
public class Connector {
    public static void main(String[] args) {
        try {
            System.setProperty("javax.net.ssl.trustStore", "./trust/certs.keystore");
            System.setProperty("javax.net.ssl.trustStorePassword", "123456");
            ConnectionString connString = new ConnectionString("mongodb://
<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin&ssl=true");
            MongoClientSettings settings = MongoClientSettings.builder()
                    .applyConnectionString(connString)
                    .applyToSslSettings(builder -> builder.enabled(true))
                    .applyToSslSettings(builder -> builder.invalidHostNameAllowed(true))
                    .build();
            MongoClient mongoClient = MongoClients.create(settings);
            MongoDatabase database = mongoClient.getDatabase("admin");
            //Ping the database. If the operation fails, an exception occurs.
            BsonDocument command = new BsonDocument("ping", new BsonInt64(1));
            Document commandResult = database.runCommand(command);
            System.out.println("Connect to database successfully");
        } catch (Exception e) {
            e.printStackTrace();
            System.out.println("Test failed");
        }
    }
}
```

Connection Without the SSL Certificate

NOTE

You do not need to download the SSL certificate because certificate verification on the server is not required.

If you connect to a cluster instance using Java, the format of code is as follows:

```
mongodb://<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin
```

Table 2-14 Parameter description

Parameter	Description
<username>	Current username.
<password>	Password for the current username
<instance_ip>	If you attempt to access the instance from an ECS, set *instance_ip* to the private IP address displayed on the **Basic Information** page of the instance to which you intend to connect.
	If you intend to access the instance through an EIP, set *instance_ip* to the EIP that has been bound to the instance.
	If multiple host addresses are required, list the addresses in the format of <instance_ip1>:<instance_port1>,<instance_ip2>:<instance_port2>....... Example: mongodb://username:*****@127.***.***.1:8635,127.***.***.2:8635/?authSource=admin
<instance_port>	Database port displayed on the **Basic Information** page. Default value: **8635**
<database_name>	Name of the database to be connected.
authSource	Authentication user database. The value is **admin**.

For details about the Java code, see the following example:

```
public class Connector {
    public static void main(String[] args) {
        try {
            ConnectionString connString = new ConnectionString("mongodb://
<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin");
            MongoClientSettings settings = MongoClientSettings.builder()
                    .applyConnectionString(connString)
                    .retryWrites(true)
                    .build();
            MongoClient mongoClient = MongoClients.create(settings);
            MongoDatabase database = mongoClient.getDatabase("admin");
            //Ping the database. If the operation fails, an exception occurs.
            BsonDocument command = new BsonDocument("ping", new BsonInt64(1));
            Document commandResult = database.runCommand(command);
            System.out.println("Connect to database successfully");
        } catch (Exception e) {
            e.printStackTrace();
            System.out.println("Test failed");
        }
    }
}
```

2.1.5.2 Python

This section describes how to use the MongoDB client in Python to connect to a cluster instance.

Prerequisites

1. To connect an ECS to an instance, the ECS must be able to communicate with the DDS instance. You can run the following command to connect to the IP address and port of the instance server to test the network connectivity.

   curl *ip:port*

   If the message **It looks like you are trying to access MongoDB over HTTP on the native driver port** is displayed, the network connectivity is normal.
2. Install Python and third-party installation package **pymongo** on the ECS. Pymongo 2.8 is recommended.
3. If SSL is enabled, you need to download the root certificate and upload it to the ECS.

Connection Code

- Enabling SSL

```
import ssl
from pymongo import MongoClient
conn_urls="mongodb://rwuser:rwuserpassword@ip:port/{mydb}?authSource=admin"
connection = MongoClient(conn_urls,connectTimeoutMS=5000,ssl=True,
ssl_cert_reqs=ssl.CERT_REQUIRED,ssl_match_hostname=False,ssl_ca_certs=${path to
certificate authority file})
dbs = connection.database_names()
print "connect database success! database names is %s" % dbs
```

- Disabling SSL

```
import ssl
from pymongo import MongoClient
conn_urls="mongodb://rwuser:rwuserpassword@ip:port/{mydb}?authSource=admin"
connection = MongoClient(conn_urls,connectTimeoutMS=5000)
dbs = connection.database_names()
print "connect database success! database names is %s" % dbs
```

NOTE

- The authentication database in the URL must be **admin**. That means setting **authSource** to **admin**.
- In SSL mode, you need to manually generate the trustStore file.
- The authentication database must be **admin**, and then switch to the service database.

2.1.5.3 PHP

This section describes how to connect to a cluster instance using PHP.

Prerequisites

1. To connect an ECS to a DDS instance, run the following command to connect to the IP address and port of the instance server to test the network connectivity.

   curl ip:port

   If the message **It looks like you are trying to access MongoDB over HTTP on the native driver port** is displayed, the ECS and DDS instance can communicate with each other.

2. If SSL is enabled, you need to download the root certificate and upload it to the ECS.

Obtaining and Using PHP

For the information about PHP, visit **https://www.php.net/mongodb-driver-manager.construct**

Connection Code

- Enabling SSL
  - Run **MongoDB\Client::__construct()** to create a client instance.

    ```
    function __construct(
        ?string $uri = null,
        array $uriOptions = [],
        array $driverOptions = []
    )
    ```

  - Use $uriOptions to set **SSL** to **true** to enable the SSL connection. Use $driverOptions to set **ca_file** to the CA certificate path and **allow_invalid_hostname** to **true**.

    ```
    <?php

    require 'vendor/autoload.php'; // include Composer goodies

    $replicaset_url = 'mongodb://rwuser:*****@192.168.***.***:8635,192.168.***.***:8635/
    test?authSource=admin';
    $test_db = 'test_db';
    $test_coll = 'test_coll';

    //Create mongoclient.
    $client = new MongoDB\Client(
    ....$replicaset_url,
        [
            'ssl' => true,
        ],
        [
            "ca_file" => "/path/to/ca.pem",
            "allow_invalid_hostname" => true
        ]
    );

    $collection = $client->$test_db->$test_coll;

    //Insert a record.
    $result = $collection->insertOne([
        'username' => 'admin',
        'email' => 'admin@example.com',
    ]);
    echo "Object ID: '{$result->getInsertedId()}'", "\n";

    //Query a record.
    $result = $collection->find(['username' => 'admin']);
    foreach ($result as $entry) {
        echo $entry->_id, ': ', $entry->email, "\n";
    }
    ```

```
?>
```

- Disabling SSL

```
<?php

require 'vendor/autoload.php'; // include Composer goodies

$replicaset_url = 'mongodb://rwuser:*****@192.168.***.***:8635,192.168.***.***:8635/test?
authSource=admin';
$test_db = 'test_db';
$test_coll = 'test_coll';

//Create mongoclient.
$client = new MongoDB\Client($replicaset_url);
$collection = $client->$test_db->$test_coll;

//Insert a record.
$result = $collection->insertOne([
    'username' => 'admin',
    'email' => 'admin@example.com',
]);
echo "Object ID: '{$result->getInsertedId()}'", "\n";

//Query a record.
$result = $collection->find(['username' => 'admin']);
foreach ($result as $entry) {
    echo $entry->_id, ': ', $entry->email, "\n";
}

?>
```

NOTE

- The authentication database in the URL must be **admin**. Set **authSource** to **admin**.
- Change the authentication database of the **rwuser** user to **admin**, and then switch to the service database after authentication.

2.2 Connecting to a Replica Set Instance

2.2.1 Connection Methods

You can access DDS over private or public networks.

Table 2-15 Connection methods

Method	IP Address	Scenario	Description
DAS	Not required	DAS provides a GUI and allows you to perform visualized operations on the console. SQL execution, advanced database management, and intelligent O&M are available to make database management simple, secure, and intelligent.	● Easy to use, secure, advanced, and intelligent ● Recommended

Method	IP Address	Scenario	Description
Private network	Private IP address	DDS provides a private IP address by default. If your applications are running on an ECS in the same region, AZ, and VPC subnet as your DDS instance, you are advised to use a private IP address to connect the ECS to your DDS instances.	Secure and excellent performance
Public network	EIP	● If your applications are running on an ECS that is in a different region from the one where the DB instance is located, use an EIP to connect the ECS to your DDS DB instances. ● If your applications are deployed on another cloud platform, EIP is recommended.	● Low security ● For faster transmission and improved security, you are advised to migrate your applications to an ECS that is in the same subnet as your DDS instance and use a private IP address to access the instance.

2.2.2 (Recommended) Connecting to Replica Set Instances Through DAS

2.2.2.1 Overview

DAS provides a GUI and allows you to perform visualized operations on the console. SQL execution, advanced database management, and intelligent O&M are available to make database management simple, secure, and intelligent. You are advised to use DAS to connect to DB instances.

This section describes how to buy a replica set instance on the management console and how to connect to the replica set instance through DAS.

Process

To purchase and connect to a replica set instance, perform the following steps:

1. **Buy a replica set instance.**
2. **Connect to the replica set instance through DAS.**

2.2.2.2 Connecting to a Replica Set Instance Through DAS

Data Admin Service (DAS) enables you to manage DB instances on a web-based console, simplifying database management and improving working efficiency. You can connect and manage instances through DAS. By default, you have the permission required for remote login. It is recommended that you use the DAS service to connect to instances. DAS is secure and convenient.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

If you want compute and network resources dedicated to your exclusive use, **enable a DeC** and **apply for DCC resources**. After enabling a DeC, you can select the DeC region and project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click **Log In** in the **Operation** column.

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

Figure 2-44 Instance management

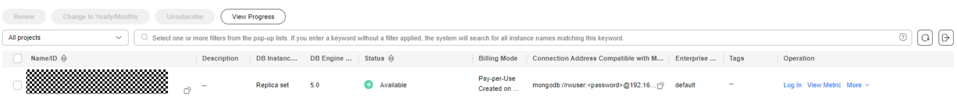
Renew
Change to Yearly/Monthly
Unsubscribe
View Progress
All projects
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all instance names matching this keyword.

Name/ID	Description	DB Instanc...	DB Engine ...	Status	Billing Mode	Connection Address Compatible with M...	Enterprise ...	Tags	Operation
[redacted]	--	Replica set	5.0	Available	Pay-per-Use Created on ...	mongodb://rwuser:<password>@192.16...	default	--	Log In View Metric More

Step 5 On the displayed login page, enter the administrator username and password and click **Log In**.

For details about how to manage databases through DAS, see **Database Management**.

----End

2.2.3 Connecting to a Replica Set Instance over a Private Network

2.2.3.1 Configuring Security Group Rules

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

To ensure database security and reliability, you need to configure security group rules to allow specific IP addresses and ports to access DDS instances.

You can connect to an instance by configuring security group rules in following two ways:

- If the ECS and instance are in the same security group, they can communicate with each other by default. No security group rule needs to be configured. Go to **Connecting to a Replica Set Instance Using Mongo Shell (Private Network)**.

Figure 2-45 Same security group

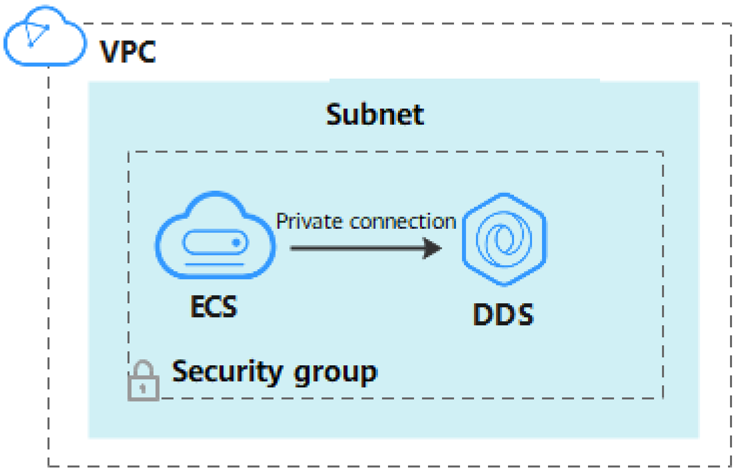

- If the ECS and instance are in different security groups, you need to configure security group rules for them, separately.

Figure 2-46 Different security groups

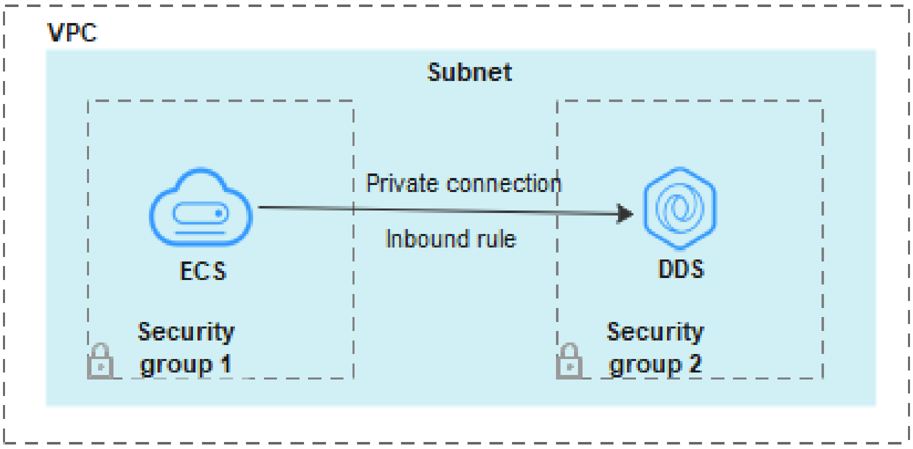

  - Instance: Configure an **inbound rule** for the security group associated with the instance.
  - ECS: The default security group rule allows all outbound data packets. In this case, you do not need to configure a security group rule for the ECS. If not all traffic is allowed to reach the instance, configure an **outbound** rule for the ECS.

This section describes how to configure an inbound rule for an instance.

Precautions

- By default, an account can create up to 500 security group rules.
- Too many security group rules will increase the first packet latency, so a maximum of 50 rules for each security group is recommended.
- By default, one DDS instance is associated with only one security group.
- DDS allows you to associate multiple security groups to a DB instance. You can apply for the service based on your service requirements. For better network performance, you are advised to select no more than five security groups.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Network Information** area on the **Basic Information** page, click the security group.

Figure 2-47 Security Group

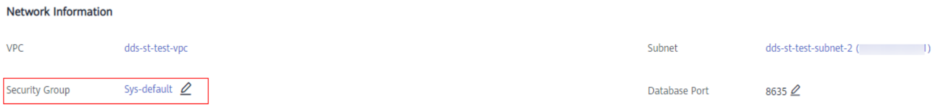

Network Information

VPC	dds-st-test-vpc	Subnet	dds-st-test-subnet-2 ([redacted])
Security Group	Sys-default	Database Port	8635

You can also choose **Connections** in the navigation pane on the left. On the **Private Connection** tab, in the **Security Group** area, click the security group name.

Figure 2-48 Security Group

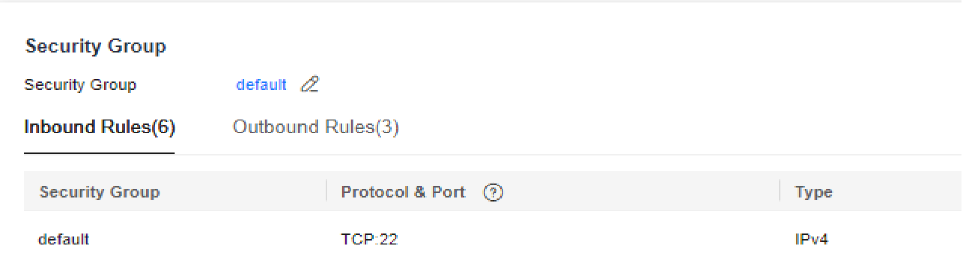

Security Group

Security Group: default

Inbound Rules(6) Outbound Rules(3)

Security Group	Protocol & Port	Type
default	TCP:22	IPv4

Step 6 On the **Security Group** page, locate the target security group and click **Manage Rule** in the **Operation** column.

Step 7 On the **Inbound Rules** tab, click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Step 8 Add a security group rule as prompted.

Figure 2-49 Add Inbound Rule

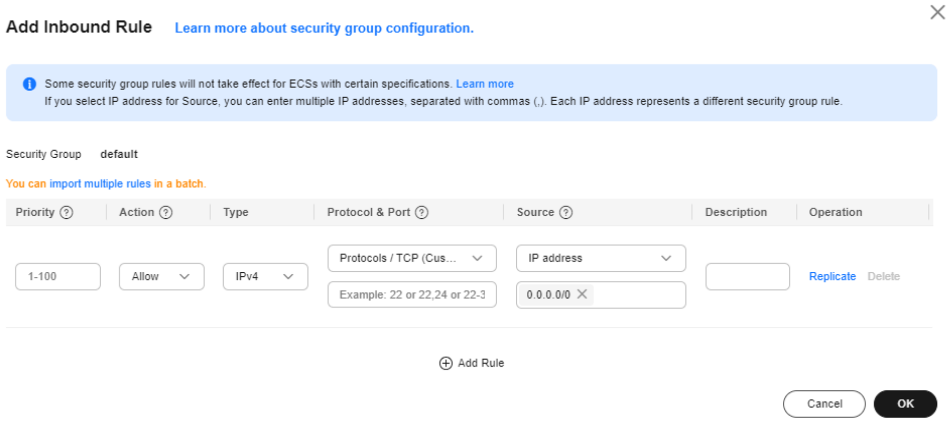

Add Inbound Rule Learn more about security group configuration.

Some security group rules will not take effect for ECSs with certain specifications. Learn more
If you select IP address for Source, you can enter multiple IP addresses, separated with commas (,). Each IP address represents a different security group rule.

Security Group default

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Source	Description	Operation
1-100	Allow	IPv4	Protocols / TCP (Cus... Example: 22 or 22,24 or 22-3	IP address 0.0.0.0/0		Replicate Delete

Table 2-16 Inbound rule settings

Parameter	Description	Example
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default priority is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	The security group rule actions. A rule with a deny action overrides another with an allow action if the two rules have the same priority.	Allow
Protocol & Port	The network protocol required for access. Available options: **TCP**, **UDP**, **ICMP**, or **GRE**	TCP
	Port: the port on which you wish to allow access to DDS. The default port is 8635. The port ranges from 2100 to 9500 or can be 27017, 27018, or 27019.	8635
Type	IP address type. Only **IPv4** and **IPv6** are supported.	IPv4

Parameter	Description	Example
Source	Specifies the supported IP address, security group, and IP address group, which allow access from IP addresses or instances in other security group. Example: ● Single IP address: 192.168.10.10/32 ● IP address segment: 192.168.1.0/24 ● All IP addresses: 0.0.0.0/0 ● Security group: sg-abc ● IP address group: ipGroup-test If you enter a security group, all ECSs associated with the security group comply with the created rule. For more information about IP address groups, see **IP Address Group Overview**.	0.0.0.0/0
Description	(Optional) Provides supplementary information about the security group rule. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

Step 9 Click **OK**.

----End

2.2.3.2 Connecting to a Replica Set Instance Using Mongo Shell (Private Network)

Mongo shell is the default client for the MongoDB database server. You can use Mongo Shell to connect to DB instances, and query, update, and manage data in databases. DDS is compatible with MongoDB. Mongo Shell is a part of the MongoDB client. To use Mongo Shell, download and install the MongoDB client first, and then use the Mongo shell to connect to the DB instance.

By default, a DDS instance provides a private IP address. If your applications are deployed on an ECS and are in the same region and VPC as DDS instances, you can connect to DDS instances using a private IP address to achieve a fast transmission rate and high security.

This section describes how to use Mongo Shell to connect to a replica set instance over a private network.

You can connect to an instance using an SSL connection or an unencrypted connection. The SSL connection is encrypted and more secure. To improve data transmission security, connect to instances using SSL.

Prerequisites

1. For details about how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS**.
2. Install the MongoDB client on the ECS. To ensure successful authentication, install the MongoDB client of the same version as the target instance.

   For details about how to install a MongoDB client, see **How Can I Install a MongoDB Client?**
3. The ECS can communicate with the DDS instance. For details, see **Configuring Security Group Rules**.

SSL Connection

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Connections**.

Step 6 In the **Basic Information** area, click ⤓ next to the **SSL** field.

Step 7 Upload the root certificate to the ECS to be connected to the instance.

The following describes how to upload the certificate to a Linux and Windows ECS:

- In Linux, run the following command:

  ```
  scp<IDENTITY_FILE><REMOTE_USER>@<REMOTE_ADDRESS>:<REMOTE_DIR>
  ```

  NOTE

  - **IDENTITY_FILE** is the directory where the root certificate resides. The file access permission is 600.
  - **REMOTE_USER** is the ECS OS user.
  - **REMOTE_ADDRESS** is the ECS address.
  - **REMOTE_DIR** is the directory of the ECS to which the root certificate is uploaded.
- In Windows, upload the root certificate using a remote connection tool.

Step 8 Connect to a DDS instance.

Method 1: Using the private HA connection address (recommended)

DDS provides the HA connection address. Using this address to connect to a replica set instance improves data read/write performance and prevents errors reported when data is written from the client after a primary/standby switchover.

Example command:

./mongo "*<Private HA connection address>***" --ssl --sslCAFile** *<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **Private HA Connection Address**: On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. Click the **Private Connection** tab and obtain the connection address of the current instance from the **Private HA Connection Address** field.

  Figure 2-50 Obtaining the private HA connection address

  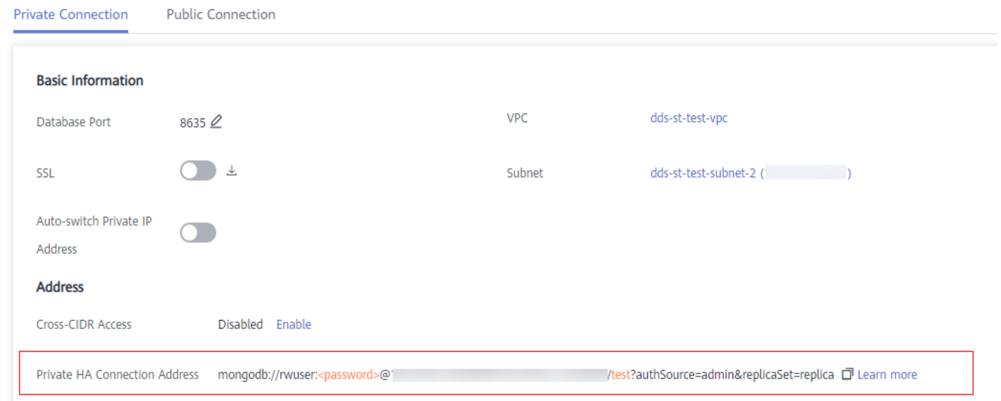

  Private Connection Public Connection

  Basic Information

  Database Port 8635

  VPC dds-st-test-vpc

  SSL

  Subnet dds-st-test-subnet-2 ([redacted])

  Auto-switch Private IP Address

  Address

  Cross-CIDR Access Disabled Enable

  Private HA Connection Address mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

  The format of the private HA connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

  mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin&replicaSet=replica**

  Pay attention to the following parameters in the private HA address:

  Table 2-17 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.

Parameter	Description
192.168.xx.xx:8635,192.168.xx.xx:8635	IP addresses and ports of the nodes of the replica set instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin&replicaSet=replica	– The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command. – **replica** in **replicaSet=replica** is the name of a replica set. The default replica set of Huawei Cloud DDS is **replica**.

- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: The replica set certificate is generated using the internal management IP address to ensure that internal communication does not occupy resources such as the user IP address and bandwidth. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

Command example:

./mongo "mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin&replicaSet=replica" --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

NOTE

- If you connect to an instance over a private HA address, add double quotation marks before and after the connection information.
- For details about the HA connection, see **Connecting to a Replica Set Instance for Read and Write Separation and High Availability**.

If the following information is displayed, the instance is successfully connected:

```
replica:PRIMARY>
```

Run the following command to access the local database:

use local

Information similar to the following is displayed:

```
switched to db local
```

Run the following command to query replica set oplog:

db.oplog.rs.find()

Method 2: Using the private HA connection address (user-defined database and account)

Example command:

./mongo "*<Private HA connection address>***" --ssl --sslCAFile** *<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **Private HA Connection Address**: On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. Click the **Private Connection** tab and obtain the connection address of the current instance from the **Private HA Connection Address** field.

Figure 2-51 Obtaining the private HA connection address

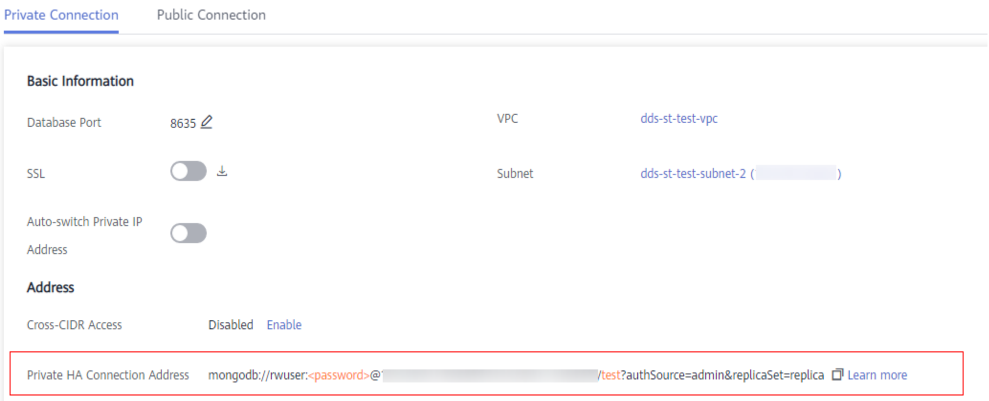

Private Connection | Public Connection

Basic Information

Database Port 8635

SSL

Auto-switch Private IP Address

VPC dds-st-test-vpc

Subnet dds-st-test-subnet-2 ([redacted])

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

The format of the obtained private HA connection address is as follows:

mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin&replicaSet=replica**

Pay attention to the following parameters in the private HA address:

Table 2-18 Parameter information

Parameter	Description
rwuser	Database username. The default value is **rwuser**. You can change the value to the username based on your service requirements.
<password>	Password for the database username. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ****@%***!$, the corresponding URL code is ****%40%25***%21%24.
192.168.xx.xx:8635,192.168.xx.xx:8635	IP addresses and ports of the nodes of the replica set instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.

Parameter	Description
authSource=admin&replicaSet=replica	– The authentication database of user **rwuser** is **admin**. – In **replica in replicaSet=replica**, **replica** indicates that the instance type is replica set and the format cannot be changed. **NOTE** If you use a user-defined database for authentication, change the authentication database in the HA connection address to the name of the user-defined database. In addition, replace **rwuser** with the username created in the user-defined database.

- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: The replica set certificate is generated using the internal management IP address to ensure that internal communication does not occupy resources such as the user IP address and bandwidth. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

For example, if you create a user-defined database **Database** and user **test1** in the database, the connection command is as follows:

./mongo "mongodb://test1:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/Database?authSource=Database&replicaSet=replica" --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

Method 3: Connect to a single node.

You can also use the private IP address of a primary or secondary node to access the replica set instance. This method affects the read/write performance when **a primary/standby switchover** occurs.

Example command:

./mongo --host *<DB_HOST>* **--port** *<DB_PORT>* **-u** *<DB_USER>* **-p --authenticationDatabase admin --ssl --sslCAFile***<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **DB_HOST** is the private IP address of the primary or secondary node of the instance to be connected.

  Primary node: You can read and write data on it.

  Secondary node: You can only read data from it.

  On the **Instances** page, click the instance to go to the **Basic Information** page. Choose **Connections**. On the **Private Connection** tab, obtain the IP address of the corresponding node.

Figure 2-52 Obtaining the IP address of a node

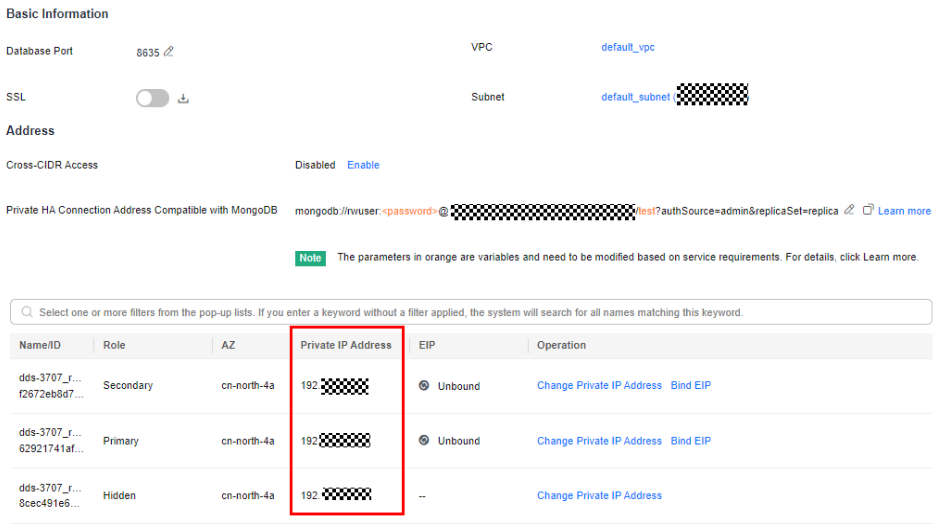

Basic Information

Database Port 8635

VPC default_vpc

SSL

Subnet default_subnet ([redacted]

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address Compatible with MongoDB mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

Note The parameters in orange are variables and need to be modified based on service requirements. For details, click Learn more.

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Role	AZ	Private IP Address	EIP	Operation
dds-3707_r... f2672eb8d7...	Secondary	cn-north-4a	192.[redacted]	Unbound	Change Private IP Address Bind EIP
dds-3707_r... 62921741af...	Primary	cn-north-4a	192[redacted]	Unbound	Change Private IP Address Bind EIP
dds-3707_r... 8cec491e6...	Hidden	cn-north-4a	192.[redacted]	--	Change Private IP Address

- **DB_PORT** is the database port. The default value is 8635.

  You can click the instance to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Private Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

  Figure 2-53 Obtaining the port

  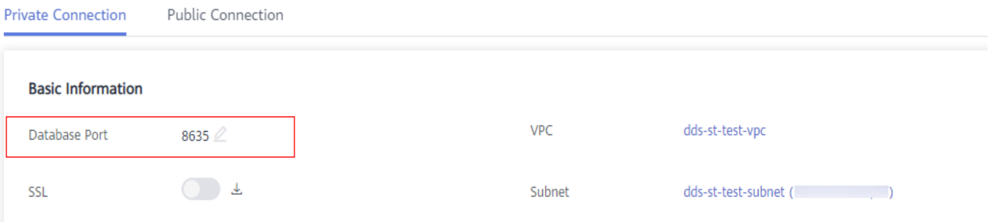

  Private Connection Public Connection

  Basic Information

  Database Port 8635

  VPC dds-st-test-vpc

  SSL

  Subnet dds-st-test-subnet ( )

- **DB_USER** is the database user. The default value is **rwuser**.
- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: The replica set certificate is generated using the internal management IP address to ensure that internal communication does not occupy resources such as the user IP address and bandwidth. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

Enter the database account password when prompted:

```
Enter password:
```

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

If the following information is displayed, the corresponding node is successfully connected:

- The primary node of the replica set is connected.

```
replica:PRIMARY>
```

- The secondary node of the replica set is connected.

```
replica:SECONDARY>
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Log in to the ECS.

Step 2 Connect to a DDS instance.

Method 1: Using the private HA connection address (recommended)

DDS provides the HA connection address. Using this address to connect to a replica set instance improves data read/write performance and prevents errors reported when data is written from the client after a primary/standby switchover.

Example command:

./mongo "*<Private HA Connection Address>***"**

Private HA Connection Address: On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. Click the **Private Connection** tab and obtain the connection address of the current instance from the **Private HA Connection Address** field.

Figure 2-54 Obtaining the private HA connection address

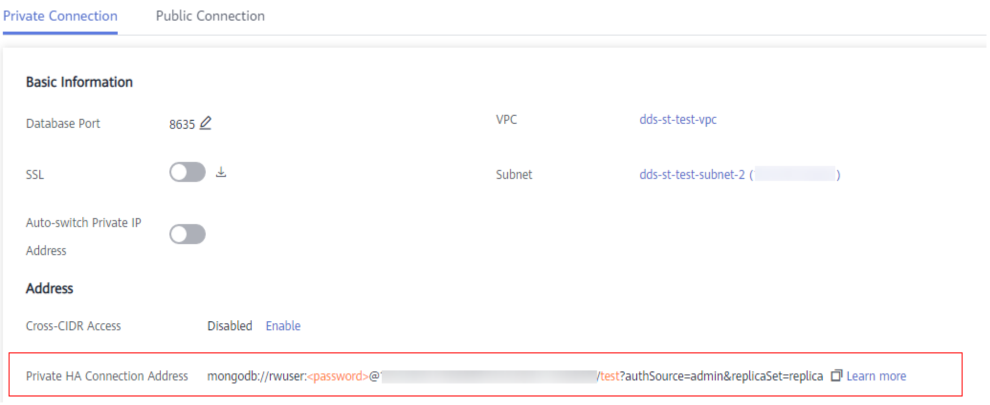

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet-2 ([redacted])

Auto-switch Private IP Address

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

The format of the private HA connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin&replicaSet=replica**

Pay attention to the following parameters in the private HA address:

Table 2-19 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ****@%***!$, the corresponding URL code is ****%40%25***%21%24.
192.168.xx.xx:8635,192.168.xx.xx:8635	IP addresses and ports of the nodes of the replica set instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin&replicaSet=replica	- The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command. - **replica** in **replicaSet=replica** is the name of a replica set. The default replica set of Huawei Cloud DDS is **replica**.

Command example:

./mongo "mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin&replicaSet=replica"**

If the following information is displayed, the instance is successfully connected:

```
replica:PRIMARY>
```

Run the following command to access the local database:

use local

Information similar to the following is displayed:

```
switched to db local
```

Run the following command to query replica set oplog:

db.oplog.rs.find()

Method 2: Using the private HA connection address (user-defined database and account)

Example command:

./mongo "*<Private HA Connection Address>***"**

Private HA Connection Address: On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. Click the **Private Connection** tab and obtain the connection address of the current instance from the **Private HA Connection Address** field.

Figure 2-55 Obtaining the private HA connection address

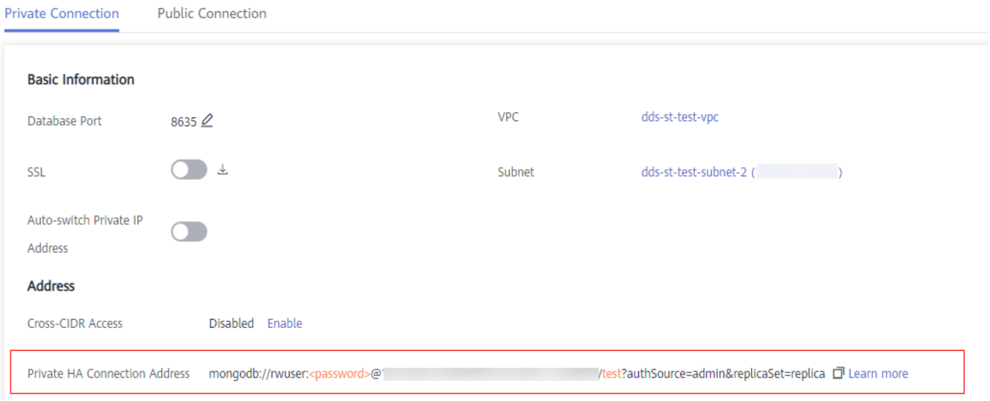

Private Connection | Public Connection

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet-2 ([redacted])

Auto-switch Private IP Address

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

The format of the obtained private HA connection address is as follows:

mongodb://rwuser:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/test?authSource=admin&replicaSet=replica**

Pay attention to the following parameters in the private HA address:

Table 2-20 Parameter information

Parameter	Description
rwuser	Database username. The default value is **rwuser**. You can change the value to the username based on your service requirements.
<password>	Password for the database username. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.xx.xx:8635,192.168.xx.xx:8635	IP addresses and ports of the nodes of the replica set instance to be connected.
test	The name of the test database. You can set this parameter based on your service requirements.

Parameter	Description
authSource=admin&replicaSet=replica	● The authentication database of user **rwuser** is **admin**. ● In **replica in replicaSet=replica**, **replica** indicates that the instance type is replica set and the format cannot be changed. **NOTE** If you use a user-defined database for authentication, change the authentication database in the HA connection address to the name of the user-defined database. In addition, replace **rwuser** with the username created in the user-defined database.

For example, if you create a user-defined database **Database** and user **test1** in the database, the connection command is as follows:

./mongo "mongodb://test1:*<password>@192.168.xx.xx:8635,192.168.xx.xx:8635***/Database?authSource=Database&replicaSet=replica"**

Method 3: Connect to a single node.

You can also use the private IP address of a primary or secondary node to access the replica set instance. This method affects the read/write performance when a primary/standby switchover occurs.

Example command:

./mongo --host *<DB_HOST>* **--port** *<DB_PORT>* **-u** *<DB_USER>* **-p --authenticationDatabase admin**

Parameter description:

- **DB_HOST** is the private IP address of the primary or secondary node of the instance to be connected.

  Primary node: You can read and write data on it.

  Secondary node: You can only read data from it.

  On the **Instances** page, click the instance to go to the **Basic Information** page. Choose **Connections**. On the **Private Connection** tab, obtain the IP address of the corresponding node.

Figure 2-56 Obtaining the IP address of a node

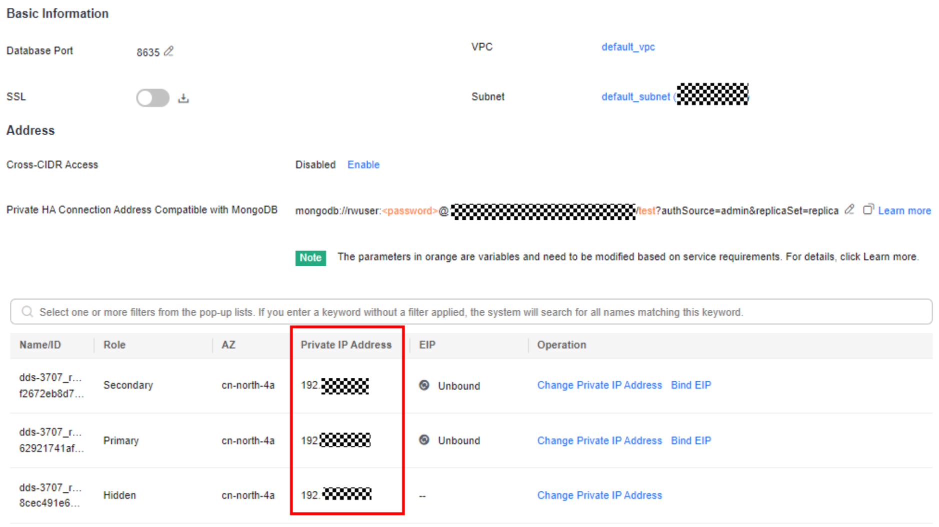

Basic Information

Database Port	8635	VPC	default_vpc
SSL		Subnet	default_subnet ([redacted]

Address

Cross-CIDR Access: Disabled Enable

Private HA Connection Address Compatible with MongoDB: mongodb://rwuser:<password>@[redacted]/test?authSource=admin&replicaSet=replica Learn more

Note The parameters in orange are variables and need to be modified based on service requirements. For details, click Learn more.

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Role	AZ	Private IP Address	EIP	Operation
dds-3707_r... f2672eb8d7...	Secondary	cn-north-4a	192.[redacted]	Unbound	Change Private IP Address Bind EIP
dds-3707_r... 62921741af...	Primary	cn-north-4a	192.[redacted]	Unbound	Change Private IP Address Bind EIP
dds-3707_r... 8cec491e6...	Hidden	cn-north-4a	192.[redacted]	--	Change Private IP Address

- **DB_PORT** is the database port. The default value is 8635.

  You can click the instance to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Private Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

Figure 2-57 Obtaining the port

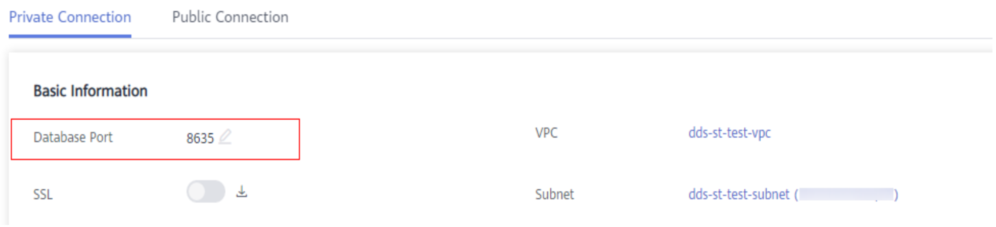

Private Connection | Public Connection

Basic Information

Database Port	8635	VPC	dds-st-test-vpc
SSL		Subnet	dds-st-test-subnet ([redacted])

- **DB_USER** is the database user. The default value is **rwuser**.

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin**

Enter the database account password when prompted:

```
Enter password:
```

If the following information is displayed, the corresponding node is successfully connected:

- The primary node of the replica set is connected.
  ```
  replica:PRIMARY>
  ```
- The secondary node of the replica set is connected.
  ```
  replica:SECONDARY>
  ```

----End

2.2.3.3 Connecting to Read Replicas Using Mongo Shell

Mongo shell is the default client for the MongoDB database server. You can use Mongo Shell to connect to DB instances, and query, update, and manage data in databases. DDS is compatible with MongoDB. Mongo Shell is a part of the MongoDB client. To use Mongo Shell, download and install the MongoDB client first, and then use the Mongo shell to connect to the DB instance.

By default, a DDS instance provides a private IP address. If your applications are deployed on an ECS and are in the same region and VPC as DDS instances, you can connect to DDS instances using a private IP address to achieve a fast transmission rate and high security.

This section describes how to use Mongo Shell to connect to a read replica over a private network.

You can connect to a read replica using an SSL connection or an unencrypted connection. The SSL connection is encrypted and more secure. To improve data transmission security, connect to instances using SSL.

Prerequisites

1. For details about how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS**.
2. Install the MongoDB client on the ECS. To ensure successful authentication, install the MongoDB client of the same version as the target instance.

   For details about how to install a MongoDB client, see **How Can I Install a MongoDB Client?**
3. The ECS can communicate with the DDS instance. For details, see **Configuring Security Group Rules**.

SSL Connection

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

Step 1 On the **Instances** page, click the instance name.

Step 2 In the navigation pane on the left, choose **Connections**.

Step 3 In the **Basic Information** area, click ⤓ next to the **SSL** field.

Step 4 Upload the root certificate to the ECS to be connected to the instance.

The following describes how to upload the certificate to a Linux and Windows ECS:

- In Linux, run the following command:

```
scp<IDENTITY_FILE><REMOTE_USER>@<REMOTE_ADDRESS>:<REMOTE_DIR>
```

NOTE

- **IDENTITY_FILE** is the directory where the root certificate resides. The file access permission is 600.
- **REMOTE_USER** is the ECS OS user.
- **REMOTE_ADDRESS** is the ECS address.
- **REMOTE_DIR** is the directory of the ECS to which the root certificate is uploaded.

- In Windows, upload the root certificate using a remote connection tool.

Step 5 Connect to a DDS instance. The DDS console provides the read replica connection address. You can use this address to connect to the read replica.

Example command:

./mongo "*<Read replica connection address>***" --ssl --sslCAFile** *<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **Read Replica Connection Address**: On the **Instances** page, click the instance to go to the **Basic Information** page. Choose **Connections**. Click the **Private Connection** tab. In the **Address** area, obtain the connection address of the read replica instance.

Figure 2-58 Obtaining the read replica connection address

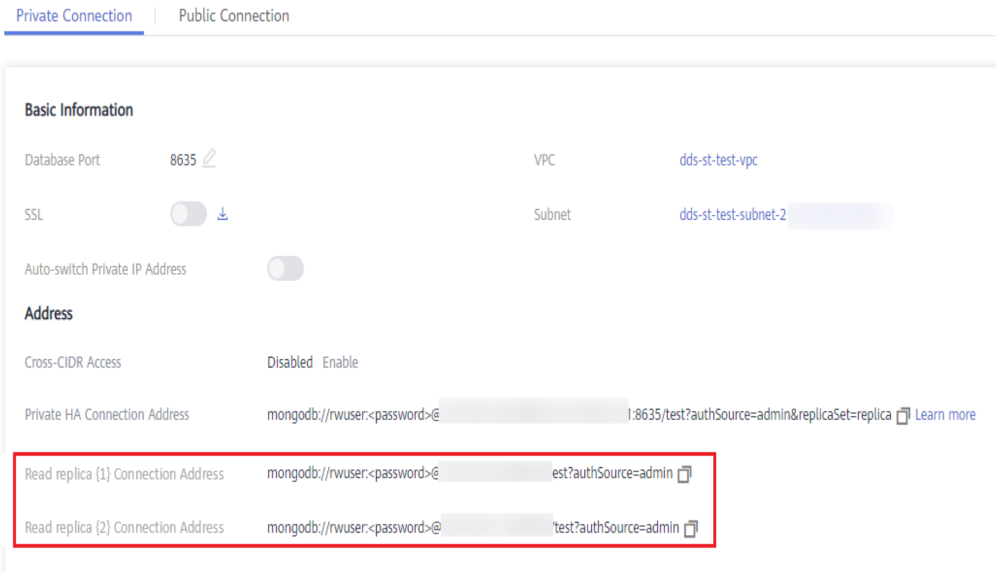

Private Connection | Public Connection

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet-2 [redacted]

Auto-switch Private IP Address

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address mongodb://rwuser:<password>@[redacted]:8635/test?authSource=admin&replicaSet=replica Learn more

Read replica {1} Connection Address mongodb://rwuser:<password>@[redacted]est?authSource=admin

Read replica {2} Connection Address mongodb://rwuser:<password>@[redacted]'test?authSource=admin

The format of the read replica connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:*<password>@192.168.xx.xx:8635***/test?authSource=admin**

Pay attention to the following parameters in the read replica connection address:

Table 2-21 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25*** %21%24**.
192.168.xx.xx:8635	IP address and port of the read replica of the replica set instance.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

- **FILE_PATH** is the path for storing the root certificate.
- --sslAllowInvalidHostnames: The replica set certificate is generated using the internal management IP address to ensure that internal communication does not occupy resources such as the user IP address and bandwidth. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

Command example:

./mongo "mongodb://rwuser:*<password>@192.168.xx.xx:8635***/test?authSource=admin" --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

NOTE

When connecting to an instance using the read replica connection address, add double quotation marks (") before and after the connection information.

If the following information is displayed, the instance is successfully connected:

```
replica:SECONDARY>
```

----End

Unencrypted Connection

> **NOTICE**
>
> If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Log in to the ECS.

Step 2 Connect to a DDS instance. The DDS console provides the read replica connection address. You can use this address to connect to the read replica.

Example command:

./mongo "*<Read replica connection address>***"**

Read Replica Connection Address: On the **Instances** page, click the instance to go to the **Basic Information** page. Choose **Connections**. Click the **Private Connection** tab. In the **Address** area, obtain the connection address of the read replica instance.

Figure 2-59 Obtaining the read replica connection address

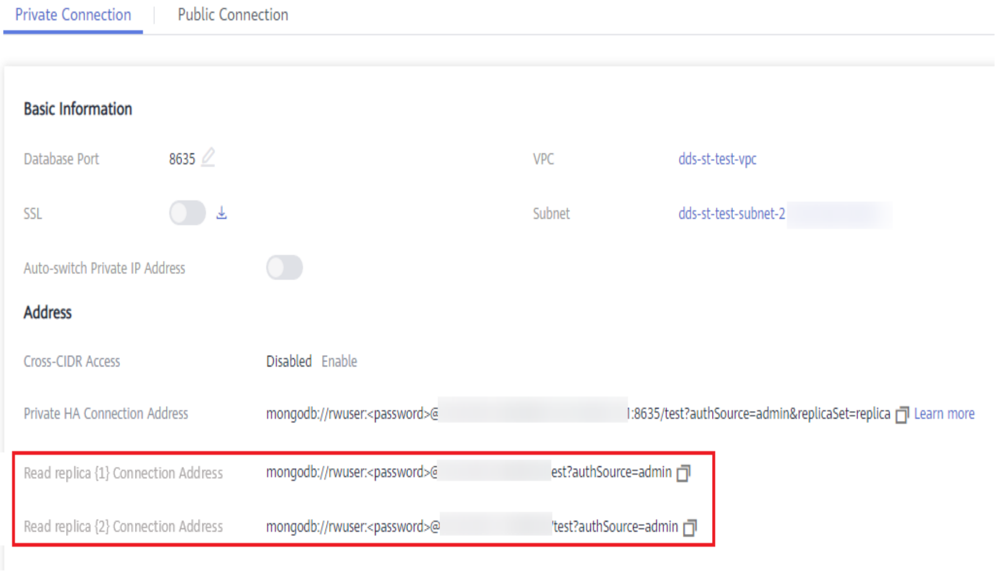

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet-2[redacted]

Auto-switch Private IP Address

Address

Cross-CIDR Access Disabled Enable

Private HA Connection Address mongodb://rwuser:<password>@[redacted]:8635/test?authSource=admin&replicaSet=replica Learn more

Read replica {1} Connection Address mongodb://rwuser:<password>@[redacted]est?authSource=admin

Read replica {2} Connection Address mongodb://rwuser:<password>@[redacted]'test?authSource=admin

The format of the read replica connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:*<password>@192.168.xx.xx:8635***/test?authSource=admin**

Pay attention to the following parameters in the private HA address:

Table 2-22 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ****@%***!$, the corresponding URL code is ****%40%25***%21%24.
192.168.xx.xx:8635	IP address and port of the read replica of the replica set instance.
test	The name of the test database. You can set this parameter based on your service requirements.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

Command example:

./mongo "mongodb://rwuser:*<password>@192.168.xx.xx:8635***/test?authSource=admin"**

If the following information is displayed, the instance is successfully connected:

```
replica:SECONDARY>
```

----End

2.2.4 Connecting to a Replica Set Instance over a Public Network

2.2.4.1 Binding an EIP

After you create an instance, you can bind an EIP to it to allow external access. If later you want to prohibit external access, you can also unbind the EIP from the DB instance.

Precautions

- Deleting a bound EIP does not mean that the EIP is unbound.
- Before accessing a database, apply for an EIP on the VPC console. Then, add an inbound rule to allow the IP addresses or IP address ranges of ECSs. For details, see **Configuring Security Group Rules**.

- In the replica set instance, only primary and secondary nodes can have an EIP bound. To change the EIP that has been bound to a node, you need to unbind it from the node first.

Binding an EIP

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance name.

Step 5 In the navigation pane on the left, choose **Connections**. Click the **Public Connection** tab. In the **Basic Information** area, locate the node you want to bind an EIP to and click **Bind EIP** in the **Operation** column.

Figure 2-60 Binding an EIP

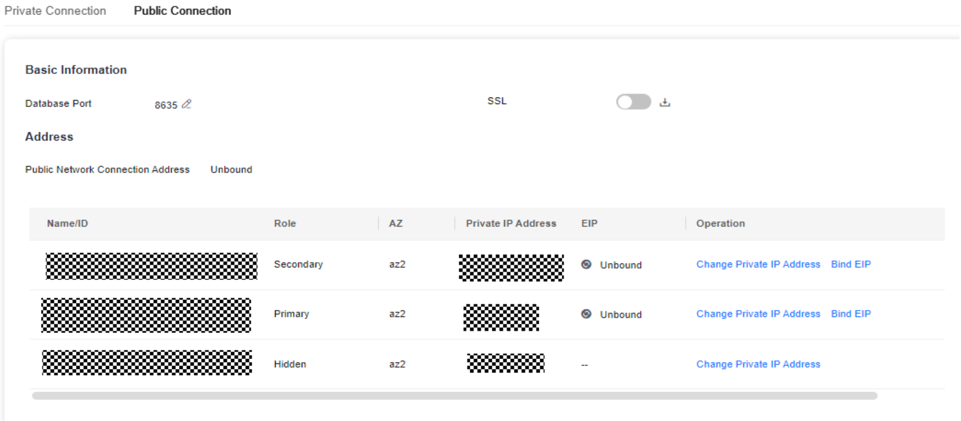

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address Unbound

Name/ID	Role	AZ	Private IP Address	EIP	Operation
[redacted]	Secondary	az2	[redacted]	Unbound	Change Private IP Address Bind EIP
[redacted]	Primary	az2	[redacted]	Unbound	Change Private IP Address Bind EIP
[redacted]	Hidden	az2	[redacted]	--	Change Private IP Address

You can also locate the node in the **Node Information area** on the **Basic Information** page and click **Bind EIP** in the **Operation** column.

Figure 2-61 Binding an EIP

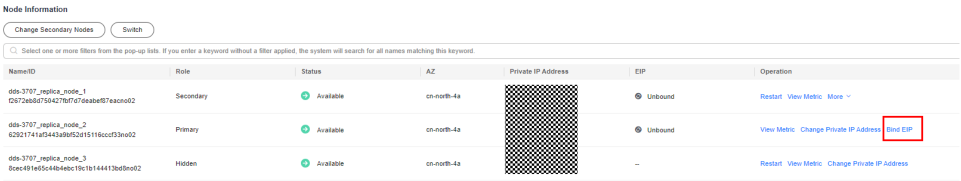

Node Information

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Role	Status	AZ	Private IP Address	EIP	Operation
dds-3707_replica_node_1 f2672eb8d750427fbf7d7deabef87eacno02	Secondary	Available	cn-north-4a	[redacted]	Unbound	Restart View Metric More
dds-3707_replica_node_2 62921741af3443a9bf52d15116cccf33no02	Primary	Available	cn-north-4a	[redacted]	Unbound	View Metric Change Private IP Address Bind EIP
dds-3707_replica_node_3 8cec491e65c44b4ebc19c1b144413bd8no02	Hidden	Available	cn-north-4a	[redacted]	--	Restart View Metric Change Private IP Address

Step 6 In the displayed dialog box, all available unbound EIPs are listed. Select the required EIP and click **OK**. If no available EIPs are displayed, click **View EIP** and create an EIP on the VPC console.

Figure 2-62 Selecting an EIP

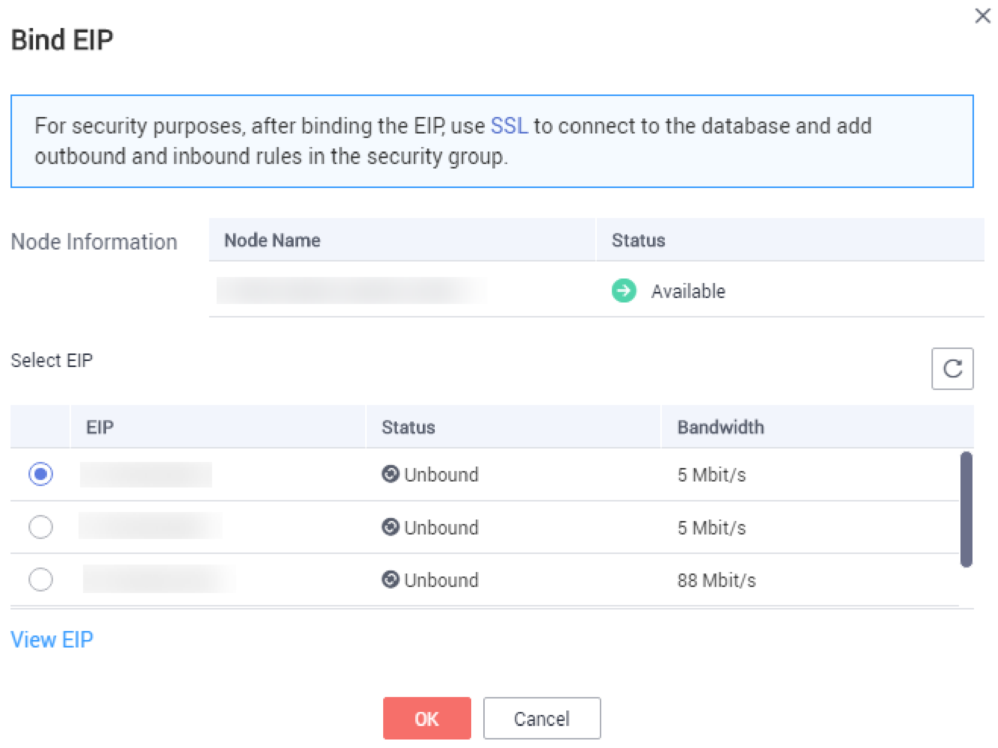

Bind EIP

For security purposes, after binding the EIP, use SSL to connect to the database and add outbound and inbound rules in the security group.

Node Information

Node Name	Status
[redacted]	Available

Select EIP

	EIP	Status	Bandwidth
●	[redacted]	Unbound	5 Mbit/s
○	[redacted]	Unbound	5 Mbit/s
○	[redacted]	Unbound	88 Mbit/s

Step 7 Locate the target node. In the **EIP** column, you can view the EIP that was bound.

To unbind an EIP from the instance, see **Unbinding an EIP**.

----End

Unbinding an EIP

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance that has been bound with an EIP.

Step 5 In the navigation pane on the left, choose **Connections**. Click the **Public Connection** tab. In the **Basic Information** area, locate the node and click **Unbind EIP** in the **Operation** column.

Figure 2-63 Unbinding an EIP

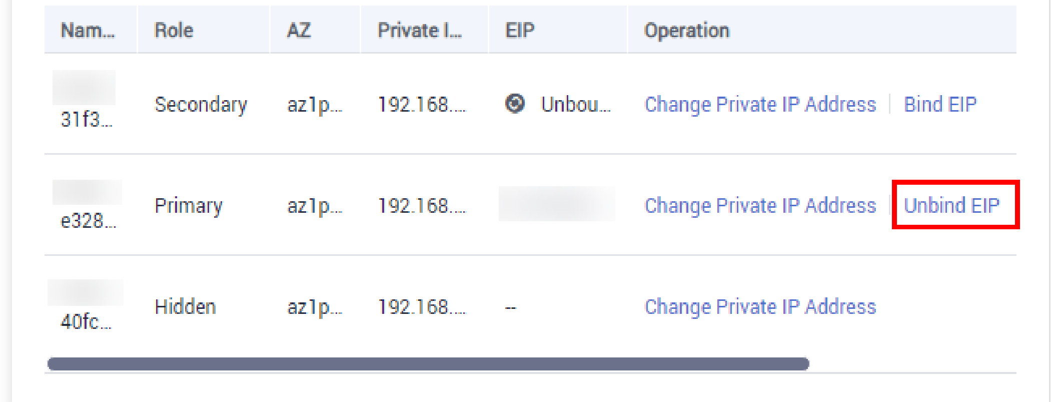

Nam...	Role	AZ	Private I...	EIP	Operation
[redacted] 31f3...	Secondary	az1p...	192.168....	⊗ Unbou...	Change Private IP Address | Bind EIP
[redacted] e328...	Primary	az1p...	192.168....	[redacted]	Change Private IP Address | Unbind EIP
[redacted] 40fc...	Hidden	az1p...	192.168....	--	Change Private IP Address

You can also locate the node in the **Node Information area** on the **Basic Information** page and click **Unbind EIP** in the **Operation** column.

Step 6 In the displayed dialog box, click **Yes**.

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

----End

2.2.4.2 Configuring Security Group Rules

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

To ensure database security and reliability, you need to configure security group rules to allow specific IP addresses and ports to access the instance.

If you attempt to connect to an instance through an EIP, you need to configure an inbound rule for the security group associated with the instance.

Precautions

- By default, an account can create up to 500 security group rules.
- Too many security group rules will increase the first packet latency, so a maximum of 50 rules for each security group is recommended.
- By default, one DDS instance is associated with only one security group.
- DDS allows you to associate multiple security groups to a DB instance. You can apply for the service based on your service requirements. For better network performance, you are advised to select no more than five security groups.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Network Information** area on the **Basic Information** page, click the security group.

Figure 2-64 Security Group

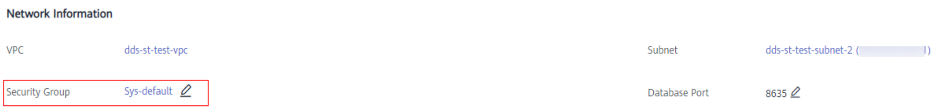

Network Information

VPC	dds-st-test-vpc	Subnet	dds-st-test-subnet-2 ([redacted])
Security Group	Sys-default	Database Port	8635

You can also choose **Connections** in the navigation pane on the left. On the **Private Connection** tab, in the **Security Group** area, click the security group name.

Figure 2-65 Security Group

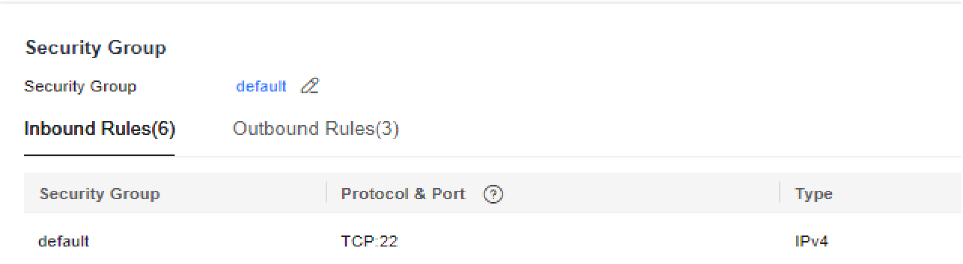

Security Group

Security Group default

Inbound Rules(6) Outbound Rules(3)

Security Group	Protocol & Port	Type
default	TCP:22	IPv4

Step 6 On the **Security Group** page, locate the target security group and click **Manage Rule** in the **Operation** column.

Step 7 On the **Inbound Rules** tab, click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Step 8 Add a security group rule as prompted.

Figure 2-66 Add Inbound Rule

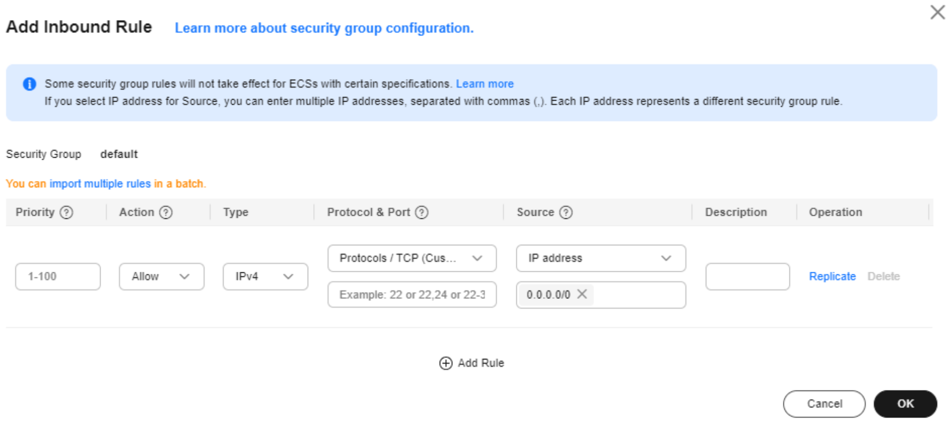

Add Inbound Rule Learn more about security group configuration.

Some security group rules will not take effect for ECSs with certain specifications. Learn more
If you select IP address for Source, you can enter multiple IP addresses, separated with commas (,). Each IP address represents a different security group rule.

Security Group default

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Source	Description	Operation
1-100	Allow	IPv4	Protocols / TCP (Cus... Example: 22 or 22,24 or 22-3	IP address 0.0.0.0/0		Replicate Delete

Table 2-23 Inbound rule settings

Parameter	Description	Example
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default priority is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	The security group rule actions. A rule with a deny action overrides another with an allow action if the two rules have the same priority.	Allow
Protocol & Port	The network protocol required for access. Available options: **TCP**, **UDP**, **ICMP**, or **GRE**	TCP
	Port: the port on which you wish to allow access to DDS. The default port is 8635. The port ranges from 2100 to 9500 or can be 27017, 27018, or 27019.	8635
Type	IP address type. Only **IPv4** and **IPv6** are supported.	IPv4
Source	Specifies the supported IP address, security group, and IP address group, which allow access from IP addresses or instances in other security group. Example: ● Single IP address: 192.168.10.10/32 ● IP address segment: 192.168.1.0/24 ● All IP addresses: 0.0.0.0/0 ● Security group: sg-abc ● IP address group: ipGroup-test If you enter a security group, all ECSs associated with the security group comply with the created rule. For more information about IP address groups, see **IP Address Group Overview**.	0.0.0.0/0

Parameter	Description	Example
Description	(Optional) Provides supplementary information about the security group rule. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

Step 9 Click **OK**.

----End

2.2.4.3 Connecting to a Replica Set Instance Using Mongo Shell (Public Network)

In the following scenarios, you can access a DDS instance from the Internet by binding an EIP to the instance.

Scenario 1: Your applications are running on an ECS that is in a different region from the one where the DDS instance is located.

Figure 2-67 Accessing DDS from ECS across regions

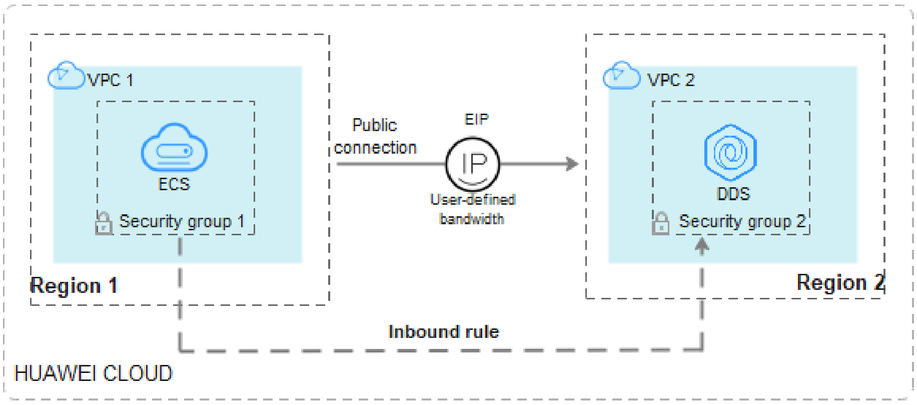

Scenario 2: Your applications are deployed on a cloud server provided by other vendors.

Figure 2-68 Accessing DDS from other cloud servers

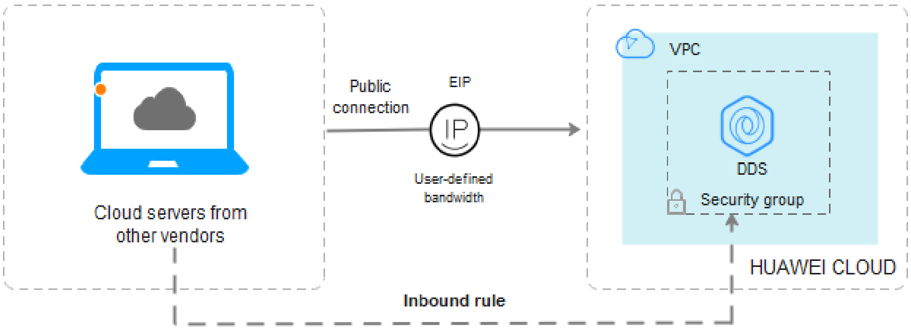

This section describes how to use Mongo Shell to connect to a replica set instance through an EIP.

You can connect to an instance using an SSL connection or an unencrypted connection. The SSL connection is encrypted and more secure. To improve data transmission security, connect to instances using SSL.

Prerequisites

1. For details about how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS**.
2. **Bind an EIP** to the replica set instance and **configure security group rules** to ensure that the replica set instance can be accessed from an ECS.
3. Install the MongoDB client on the ECS.

   For details about how to install a MongoDB client, see **How Can I Install a MongoDB Client?**

   NOTE

   The version of the installed MongoDB client must be the same as the instance version.

SSL Connection

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Connections**.

Step 6 In the **Basic Information** area, click ⤓ next to the **SSL** field.

Step 7 Upload the root certificate to the ECS to be connected to the instance.

The following describes how to upload the certificate to a Linux and Windows ECS:

- In Linux, run the following command:

  scp*<IDENTITY_FILE><REMOTE_USER>***@***<REMOTE_ADDRESS>***:***<REMOTE_DIR>*

  NOTE

  - **IDENTITY_FILE** is the directory where the root certificate resides. The file access permission is 600.
  - **REMOTE_USER** is the ECS OS user.
  - **REMOTE_ADDRESS** is the ECS address.
  - **REMOTE_DIR** is the directory of the ECS to which the root certificate is uploaded.
- In Windows, upload the root certificate using a remote connection tool.

Step 8 Connect to the instance in the directory where the MongoDB client is located.

Method 1: Using a public network connection address

Example command:

./mongo "*<Public network connection address>***" --ssl --sslCAFile***<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **Public Network Connection Address**: On the **Instances** page, click the instance to switch to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. Click the **Public Connection** tab and obtain the public network connection address.

  Figure 2-69 Obtaining the public network connection address

  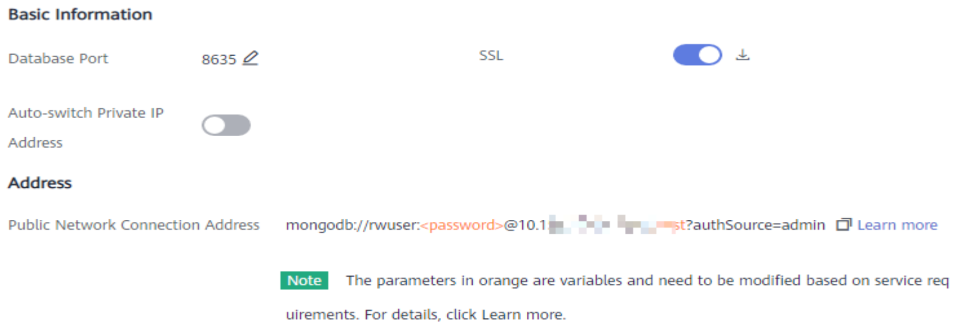

  Basic Information

  Database Port 8635

  SSL

  Auto-switch Private IP Address

  Address

  Public Network Connection Address mongodb://rwuser:<password>@10.1[redacted]st?authSource=admin Learn more

  Note The parameters in orange are variables and need to be modified based on service requirements. For details, click Learn more.

  The format of the public connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

  mongodb://rwuser:*<password>@192.168.xx.xx:8635***/test?authSource=admin**

  Pay attention to the following parameters in the public connection address:

Table 2-24 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.xx.xx:8635	The EIP and port bound to the node of the replica set instance.
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: The replica set certificate is generated using the internal management IP address to ensure that internal communication does not occupy resources such as the user IP address and bandwidth. **--sslAllowInvalidHostnames** is needed for the SSL connection through a public network.

Command example:

./mongo "mongodb://rwuser:*<password>@192.168.xx.xx:8635***/test?authSource=admin" --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

NOTE

- If you connect to an instance over a public HA address, add double quotation marks before and after the connection information.
- To improve read and write performance and prevent errors from being reported when data is written from the client after a primary/standby switchover. For details about how to connect to an instance in HA mode, see **Connecting to a Replica Set Instance for Read and Write Separation and High Availability**.

Method 2: Using an EIP

Example command:

./mongo --host *<DB_HOST>* **--port** *<DB_PORT>* **-u** *<DB_USER>* **-p --authenticationDatabaseadmin --ssl --sslCAFile***<FILE_PATH>* **--sslAllowInvalidHostnames**

Parameter description:

- **DB_HOST** is the EIP bound to the instance node to be connected.

On the **Instances** page, click the instance to go to the **Basic Information** page. Choose **Connections** > **Public Connection** and obtain the EIP of the corresponding node.

- **DB_PORT** is the database port. The default port number is 8635.

  You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Public Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

  Figure 2-70 Obtaining the port

  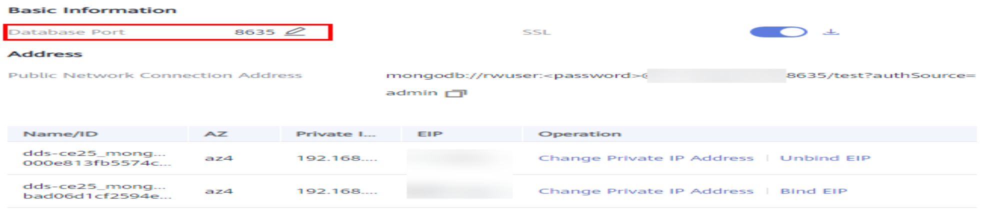

  Basic Information

  Database Port 8635

  SSL

  Address

  Public Network Connection Address mongodb://rwuser:<password>@[redacted]8635/test?authSource=admin

Name/ID	AZ	Private I...	EIP	Operation
dds-ce25_mong... 000e813fb5574c...	az4	192.168....	[redacted]	Change Private IP Address | Unbind EIP
dds-ce25_mong... bad06d1cf2594e...	az4	192.168....	[redacted]	Change Private IP Address | Bind EIP

- **DB_USER** is the database user. The default value is **rwuser**.
- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: The replica set certificate is generated using the internal management IP address to ensure that internal communication does not occupy resources such as the user IP address and bandwidth. **--sslAllowInvalidHostnames** is needed for the SSL connection through a public network.

Enter the database account password when prompted:

```
Enter password:
```

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

Step 9 Check the connection result. If the following information is displayed, the connection is successful.

- The primary node of the replica set is connected.
  ```
  replica:PRIMARY>
  ```
- The secondary node of the replica set is connected.
  ```
  replica:SECONDARY>
  ```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Log in to the ECS.

Step 2 Connect to a DDS instance.

Method 1: Using a public network connection address

Example command:

./mongo "*<Public network address>***"**

Public Network Connection Address: On the **Instances** page, click the instance to switch to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. Click the **Public Connection** tab and obtain the public network connection address.

Figure 2-71 Obtaining the public network connection address

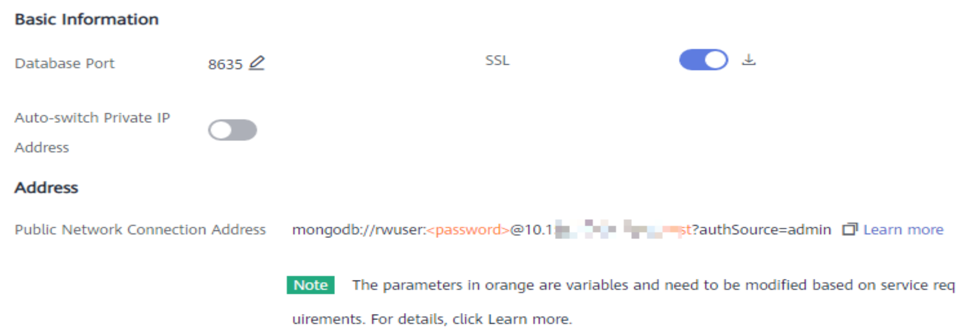

Basic Information

Database Port 8635

SSL

Auto-switch Private IP Address

Address

Public Network Connection Address mongodb://rwuser:<password>@10.1[redacted]t?authSource=admin

The parameters in orange are variables and need to be modified based on service requirements. For details, click Learn more.

The format of the public connection address is as follows. The database username **rwuser** and authentication database **admin** cannot be changed.

mongodb://rwuser:*<password>@192.168.xx.xx:8635***/test?authSource=admin**

Pay attention to the following parameters in the public connection address:

Table 2-25 Parameter description

Parameter	Description
rwuser	Account name, that is, the database username.
<password>	Password for the database account. Replace it with the actual password. If the password contains at signs (@), exclamation marks (!), dollar signs ($), or percent signs (%), replace them with hexadecimal URL codes (ASCII) %40, %21, %24, and %25 respectively. For example, if the password is ******@%***!$**, the corresponding URL code is ******%40%25***%21%24**.
192.168.xx.xx:8635	The EIP and port bound to the node of the replica set instance.

Parameter	Description
authSource=admin	The authentication database of user **rwuser** must be **admin**. **authSource=admin** is fixed in the command.

Command example:

./mongo "mongodb://rwuser:*<password>@192.168.xx.xx:8635***/test?authSource=admin"**

NOTE

- If you connect to an instance over a public HA address, add double quotation marks before and after the connection information.
- To improve read and write performance and prevent errors from being reported when data is written from the client after a primary/standby switchover. For details about how to connect to an instance in HA mode, see **Connecting to a Replica Set Instance for Read and Write Separation and High Availability**.

Method 2: Using an EIP

Example command:

./mongo --host *<DB_HOST>* **--port** *<DB_PORT>* **-u** *<DB_USER>* **-p --authenticationDatabase admin**

Parameter description:

- **DB_HOST** is the EIP bound to the instance node to be connected.

  On the **Instances** page, click the instance to go to the **Basic Information** page. Choose **Connections** > **Public Connection** and obtain the EIP of the corresponding node.
- **DB_PORT** is the database port. The default port number is 8635.

  You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Public Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

  Figure 2-72 Obtaining the port

  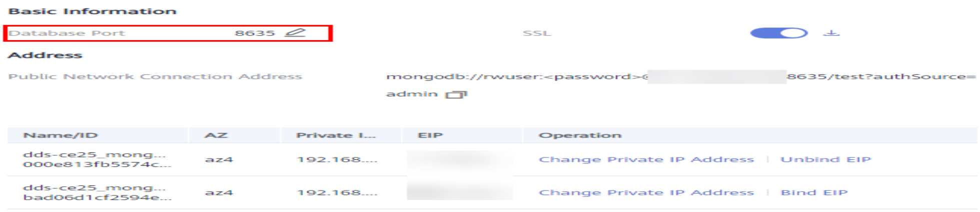

  Basic Information

  Database Port 8635

  SSL

  Address

  Public Network Connection Address mongodb://rwuser:<password>@[redacted]8635/test?authSource=admin

Name/ID	AZ	Private I...	EIP	Operation
dds-ce25_mong... 000e813fb5574c...	az4	192.168....	[redacted]	Change Private IP Address | Unbind EIP
dds-ce25_mong... bad06d1cf2594e...	az4	192.168....	[redacted]	Change Private IP Address | Bind EIP

- **DB_USER** is the database user. The default value is **rwuser**.

Enter the database account password when prompted:

```
Enter password:
```

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin**

Step 3 Check the connection result. If the following information is displayed, the connection is successful.

- The primary node of the replica set is connected.

```
replica:PRIMARY>
```

- The secondary node of the replica set is connected.

```
replica:SECONDARY>
```

----End

2.2.4.4 Connecting to a Replica Set Instance Using Robo 3T

To connect to an instance from a local device, you can use Robo 3T to access the instance from the Internet.

This section describes how to use Robo 3T to connect to a replica set instance from a local device. In this section, the Windows operating system (OS) used by the client is used as an example.

Robo 3T can connect to an instance with an unencrypted connection or an encrypted connection (SSL). To improve data transmission security, connect to instances using SSL.

Connection Diagram

Figure 2-73 Connection diagram

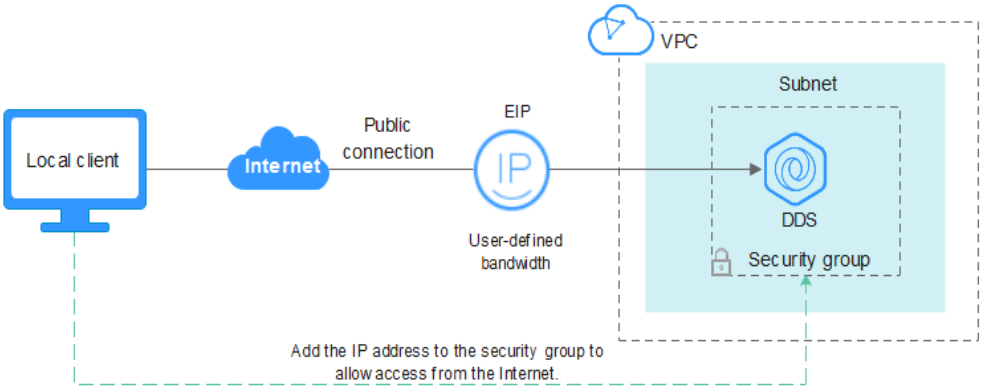

Prerequisites

1. Bind an EIP to the ECS and configure security group rules.
   a. Bind an EIP to the replica set instance.
      For details about how to bind an EIP, see **Binding an EIP**.
   b. Obtain the IP address of a local device.
   c. Configure security group rules.
      Add the IP address obtained in **1.b** and the instance port to the inbound rule of the security group.
      For details about how to configure security group rules, see **Configuring Security Group Rules**.

d. Run the ping command to ping the EIP bound in **1.a** to ensure that the EIP is accessible through your local device.

2. Install Robo 3T.

   a. For details, see **Installing Robo 3T**.

SSL

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

Step 1 Run the installed Robo 3T. On the displayed dialog box, click **Create**.

Figure 2-74 Connections

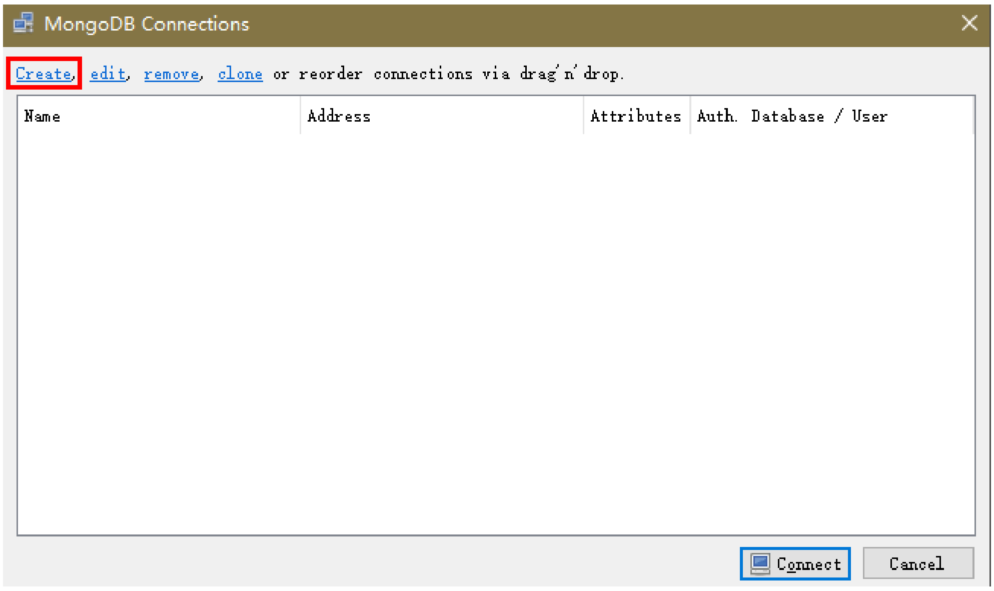

Step 2 In the **Connection Settings** dialog box, set the parameters of the new connection.

1. On the **Connection** tab, enter the name of the new connection in the **Name** text box and enter the EIP and database port that are bound to the DDS DB instance in the **Address** text box.

Figure 2-75 Connection

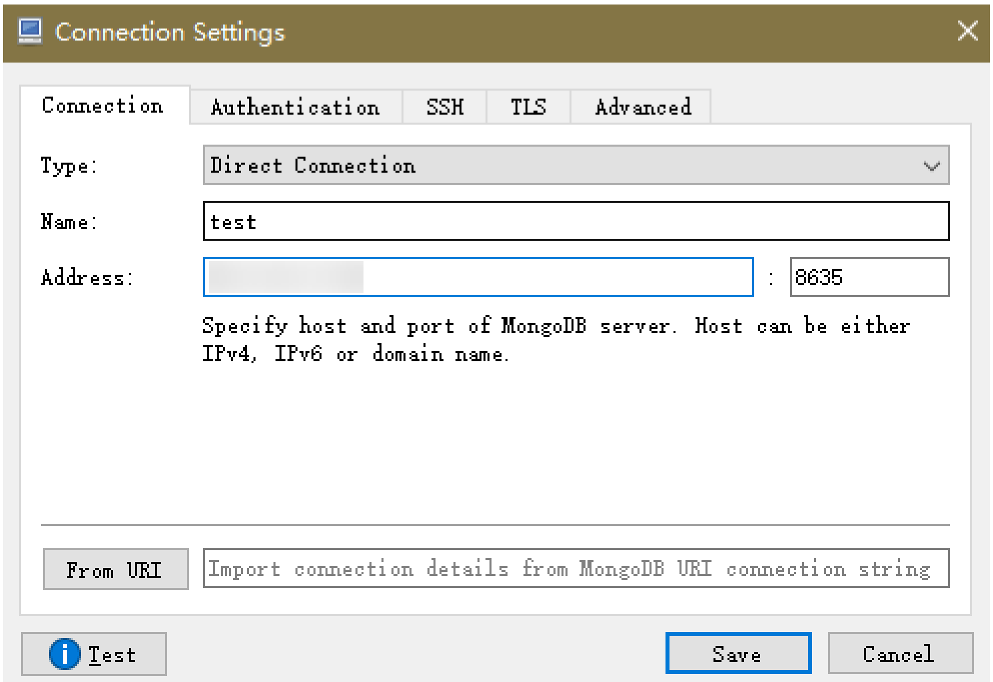

Connection Settings

Type: Direct Connection

Name: test

Address: [redacted] : 8635

Specify host and port of MongoDB server. Host can be either IPv4, IPv6 or domain name.

From URI: Import connection details from MongoDB URI connection string

2. On the **Authentication** tab, set **Database** to **admin**, **User Name** to **rwuser**, and **Password** to the administrator password you set during the creation of the cluster instance.

Figure 2-76 Authentication

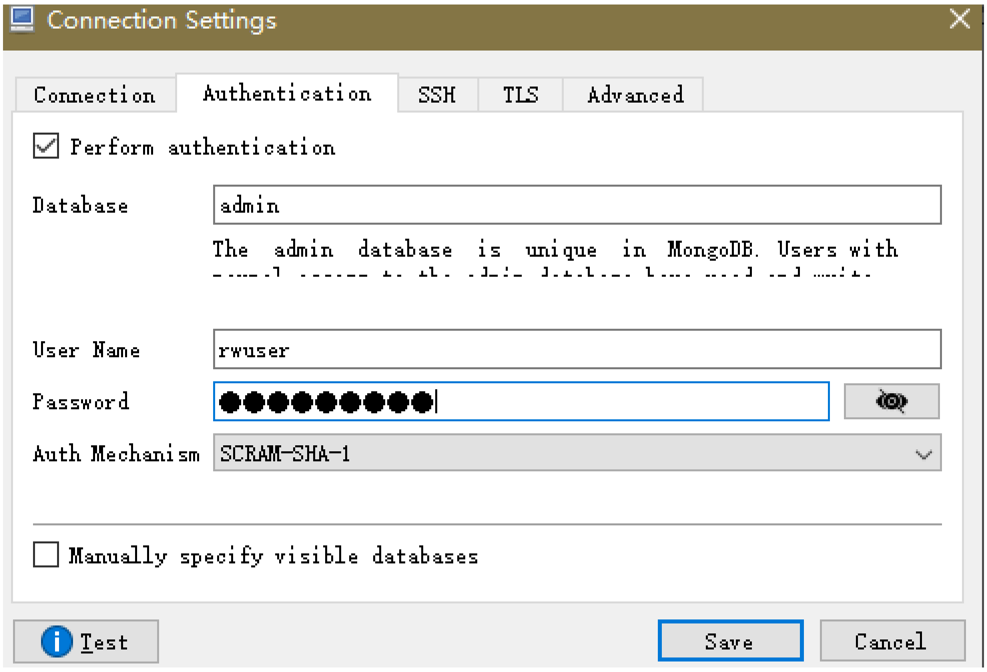

Connection Settings

☑ Perform authentication

Database: admin

The admin database is unique in MongoDB. Users with

User Name: rwuser

Password: ●●●●●●●●●●

Auth Mechanism: SCRAM-SHA-1

□ Manually specify visible databases

3. On the **TLS** tab, select **Use TLS protocol** and select **Self-signed Certificate** for **Authentication Method**.

Figure 2-77 SSL

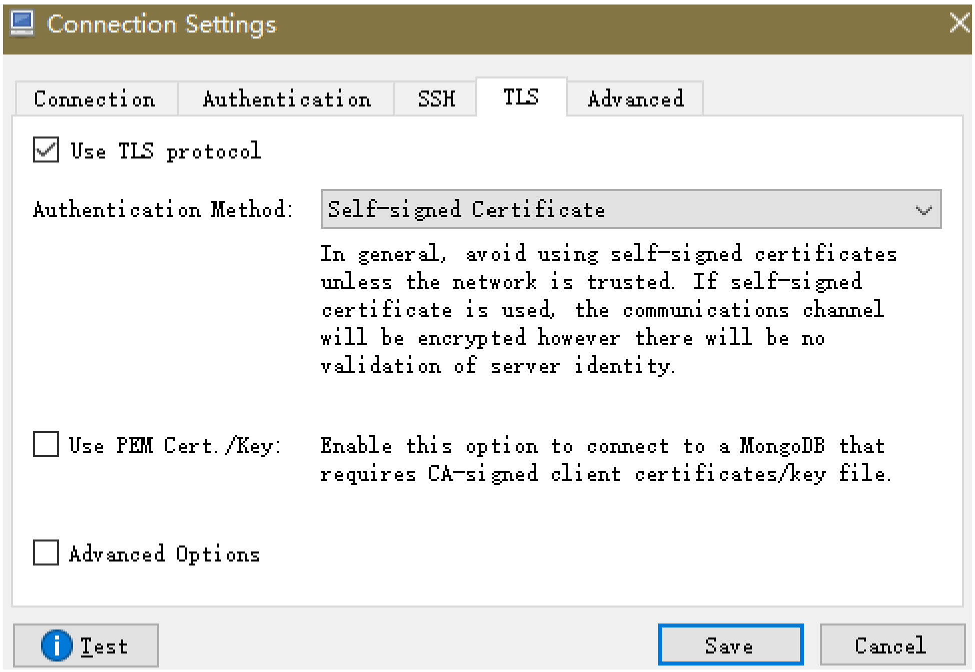

☑ Use TLS protocol

Authentication Method: Self-signed Certificate

In general, avoid using self-signed certificates unless the network is trusted. If self-signed certificate is used, the communications channel will be encrypted however there will be no validation of server identity.

□ Use PEM Cert./Key: Enable this option to connect to a MongoDB that requires CA-signed client certificates/key file.

□ Advanced Options

4. Click **Save**.

Step 3 On the **MongoDB Connections** page, click **Connect** to connect to the replica set instance.

Figure 2-78 Replica set connection information

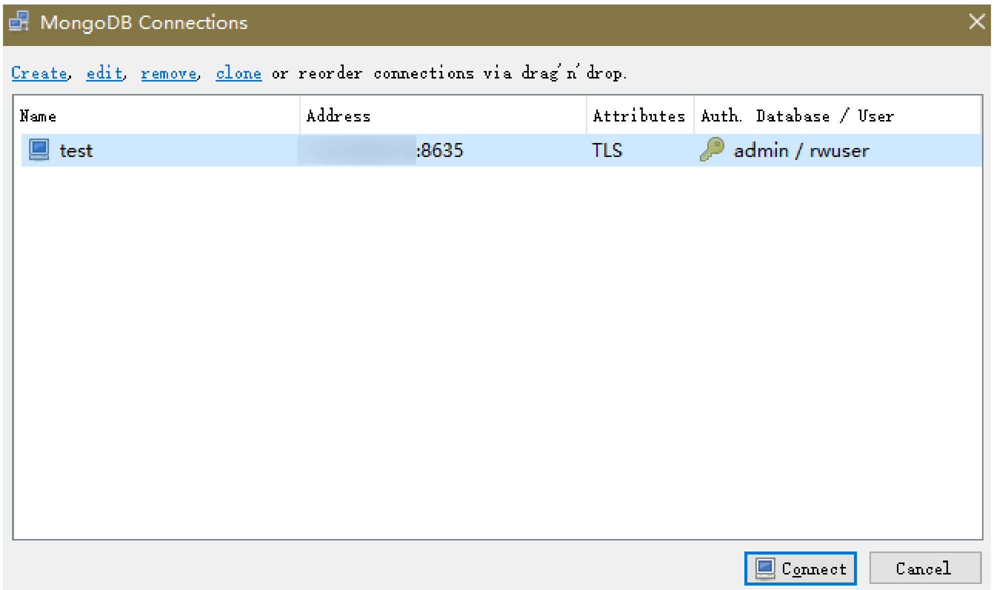

Create, edit, remove, clone or reorder connections via drag'n'drop.

Name	Address	Attributes	Auth. Database / User
test	[redacted]:8635	TLS	admin / rwuser

Step 4 If the replica set instance is successfully connected, the page shown in **Figure 2-79** is displayed.

Figure 2-79 Connection succeeded

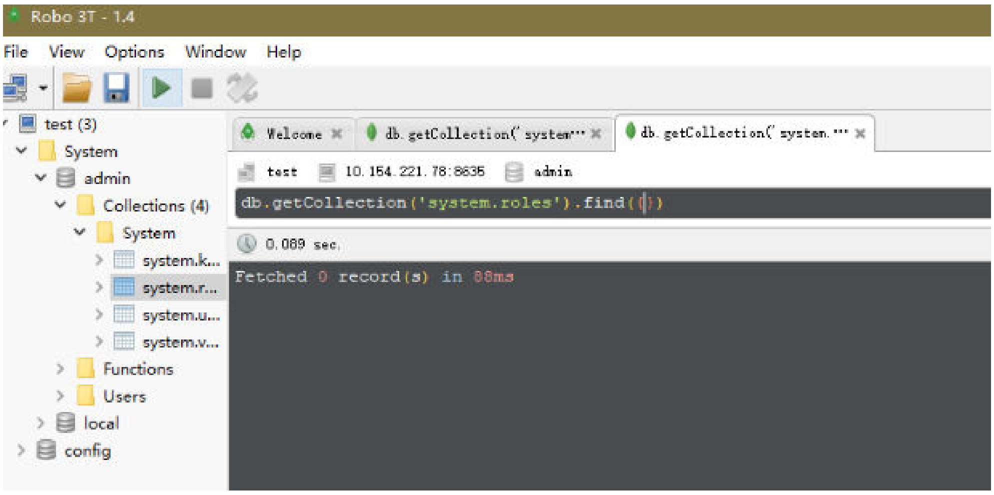

```
db.getCollection('system.roles').find({})
```

0.089 sec.

```
Fetched 0 record(s) in 88ms
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details, see **Enabling and Disabling SSL**.

Step 1 Run the installed Robo 3T. On the displayed dialog box, click **Create**.

Figure 2-80 Connections

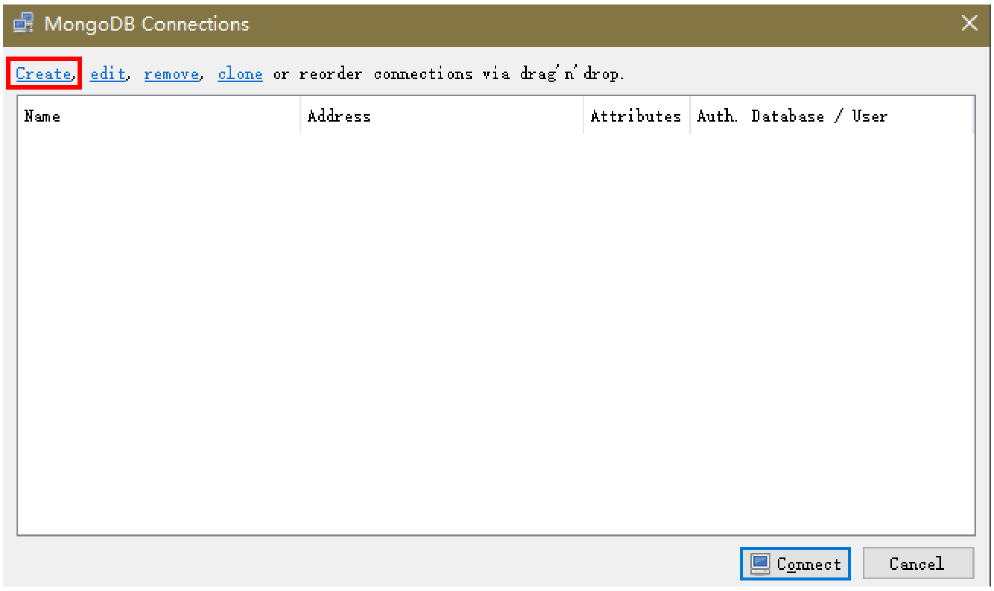

Create, edit, remove, clone or reorder connections via drag'n' drop.

Name	Address	Attributes	Auth. Database / User

Step 2 In the **Connection Settings** dialog box, set the parameters of the new connection.

1. On the **Connection** tab, enter the name of the new connection in the **Name** text box and enter the EIP and database port that are bound to the DDS DB instance in the **Address** text box.

Figure 2-81 Connection

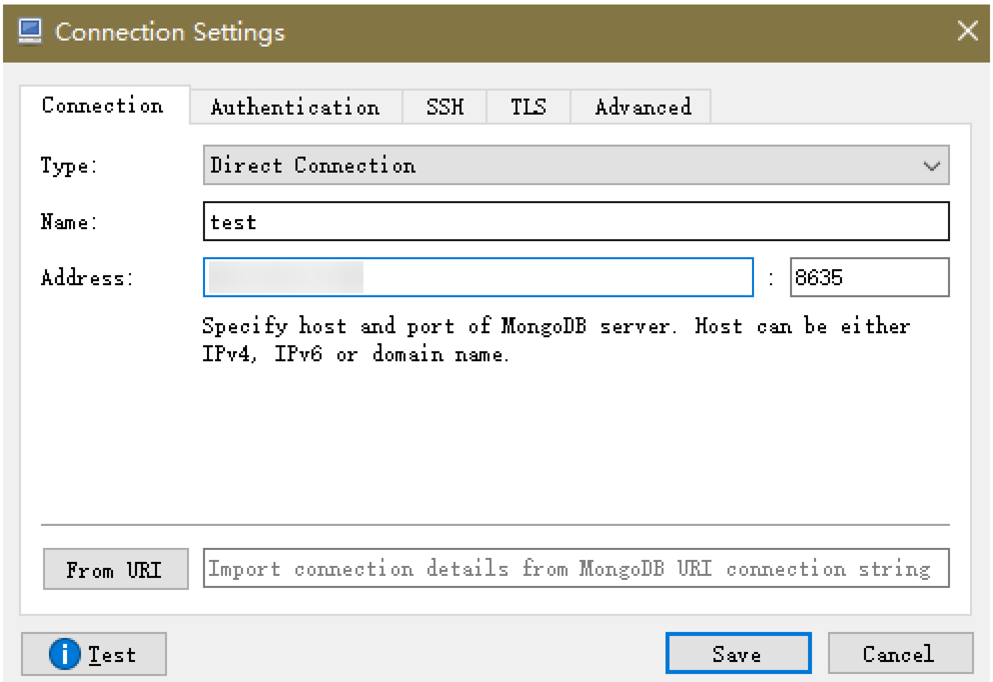

Connection Settings

Connection
Authentication
SSH
TLS
Advanced

Type: Direct Connection

Name: test

Address: [redacted] : 8635

Specify host and port of MongoDB server. Host can be either IPv4, IPv6 or domain name.

From URI
Import connection details from MongoDB URI connection string
Test
Save
Cancel

2. On the **Authentication** tab, set **Database** to **admin**, **User Name** to **rwuser**, and **Password** to the administrator password you set during the creation of the cluster instance.

Figure 2-82 Authentication

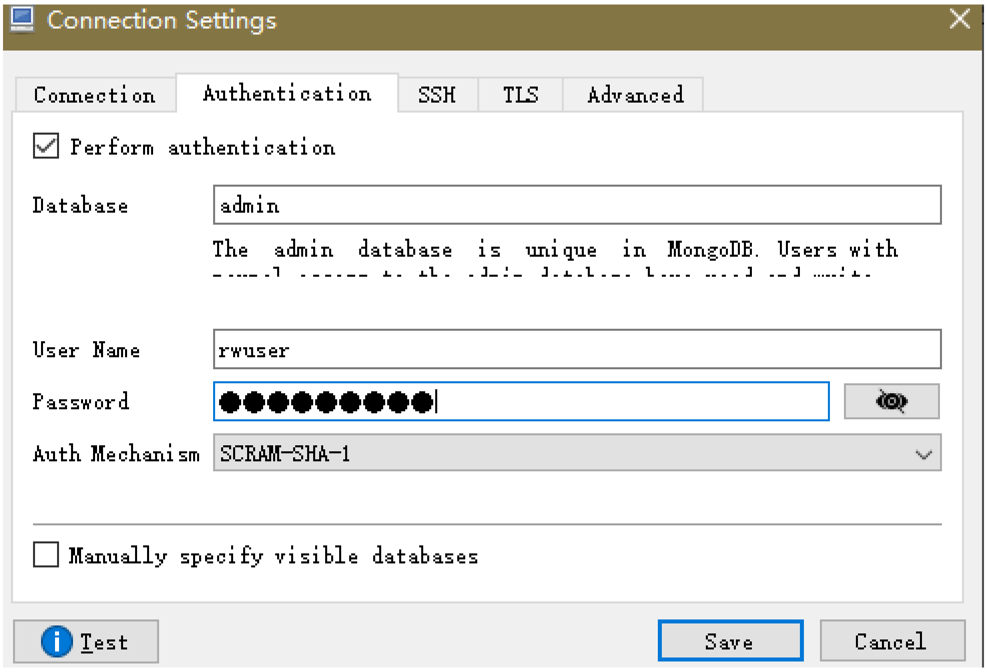

Connection Settings

☑ Perform authentication

Database: admin

The admin database is unique in MongoDB. Users with

User Name: rwuser

Password: ●●●●●●●●●●

Auth Mechanism: SCRAM-SHA-1

□ Manually specify visible databases

3. Click **Save**.

Step 3 On the **MongoDB Connections** page, click **Connect** to connect to the replica set instance.

Figure 2-83 Replica set connection information

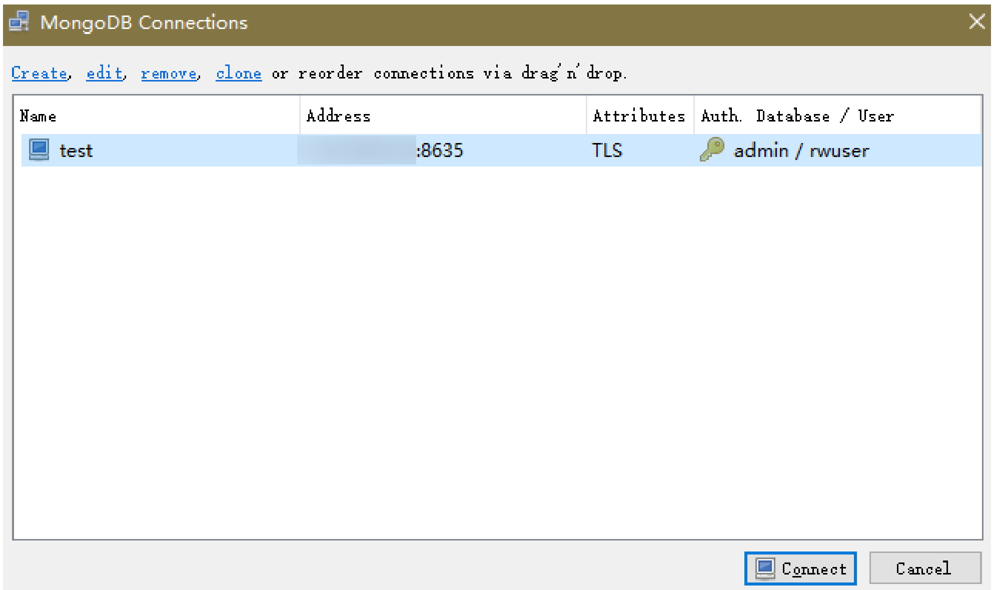

Create, edit, remove, clone or reorder connections via drag'n'drop.

Name	Address	Attributes	Auth. Database / User
test	[redacted]:8635	TLS	admin / rwuser

Step 4 If the replica set instance is successfully connected, the page shown in **Figure 2-84** is displayed.

Figure 2-84 Connection succeeded

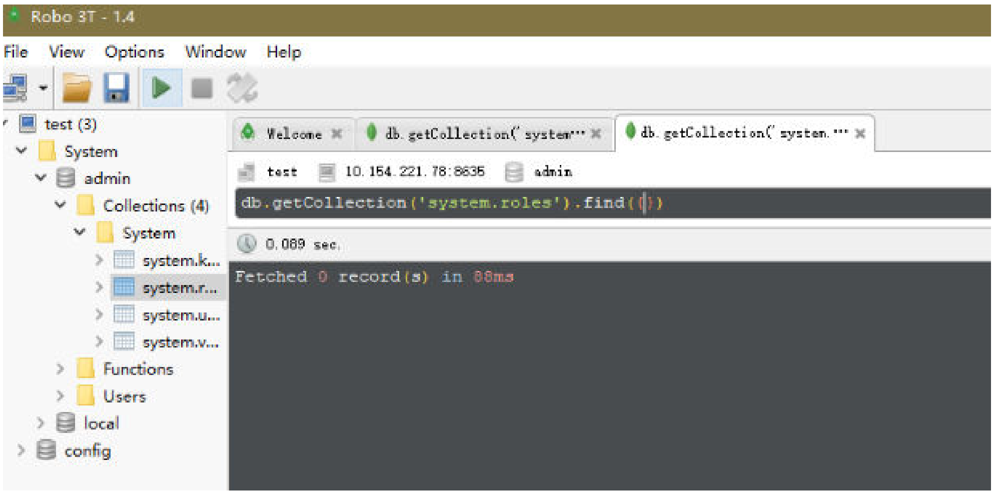

test 10.154.221.78:8635 admin

```
db.getCollection('system.roles').find({})
```

0.089 sec.

```
Fetched 0 record(s) in 88ms
```

----End

2.2.5 Connecting to a Replica Set Instance Using Program Code

2.2.5.1 Java

If you are connecting to an instance using Java, an SSL certificate is optional, but downloading an SSL certificate and encrypting the connection will improve the security of your instance. SSL is disabled by default for newly created instances, but you can enable SSL by referring to **Enabling or Disabling SSL**. SSL encrypts connections to databases but it increases the connection response time and CPU usage. For this reason, enabling SSL is not recommended.

Prerequisites

Familiarize yourself with:

- Computer basics
- Java code

Obtaining and Using Java

- Download the Jar driver from: **https://repo1.maven.org/maven2/org/mongodb/mongo-java-driver/3.0.4/**
- To view the usage guide, visit **https://mongodb.github.io/mongo-java-driver/4.2/driver/getting-started/installation/**.

Using an SSL Certificate

NOTE

- Download the SSL certificate and verify the certificate before connecting to databases.
- On the **Instances** page, click the target DB instance name. In the **DB Information** area on the **Basic Information** page, click ⤓ in the **SSL** field to download the root certificate or certificate bundle.
- For details about how to set up an SSL connection, see the MongoDB Java Driver official document at **https://www.mongodb.com/docs/drivers/java/sync/current/fundamentals/connection/tls/#std-label-tls-ssl**.
- Java Runtime Environment (JRE) earlier than Java 8 enables TLS 1.2 only in updated versions. If TLS 1.2 is not enabled for your JRE, upgrade it to a later version to use TLS 1.2 for connection.

Use Java to connect to the replica set. The format of the Java code is as follows:

```
mongodb://<username>:<password>@<instance_ip>:<instance_port>/<database_name>?authSource=admin&replicaSet=replica&ssl=true
```

Table 2-26 Parameter description

Parameter	Description
<username>	Current username.
<password>	Password for the current username
<instance_ip>	If you attempt to access the instance from an ECS, set *instance_ip* to the private IP address displayed on the **Basic Information** page of the instance to which you intend to connect.
	If you intend to access the instance through an EIP, set *instance_ip* to the EIP that has been bound to the instance.
<instance_port>	Database port displayed on the **Basic Information** page. Default value: **8635**
<database_name>	Name of the database to be connected.
authSource	Authentication user database. The value is **admin**.

Parameter	Description
ssl	Connection mode. **true** indicates that the SSL connection mode is used.

Use the keytool to configure the CA certificate. For details about the parameters, see **Table 2-27**.

```
keytool -importcert -trustcacerts -file <path to certificate authority file> -keystore <path to trust store> -
storepass <password>
```

Table 2-27 Parameter description

Parameter	Description
<path to certificate authority file>	Path for storing the SSL certificate.
<path to trust store>	Path for storing the truststore. Set this parameter as required, for example, **./trust/certs.keystore**.
<password>	Custom password.

Set the JVM system properties in the program to point to the correct truststore and keystore:

- System.setProperty("javax.net.ssl.trustStore","<path to trust store>");
- System.setProperty("javax.net.ssl.trustStorePassword","<password>");

For details about the Java code, see the following example:

```
public class Connector {
    public static void main(String[] args) {
        try {
            System.setProperty("javax.net.ssl.trustStore", "./trust/certs.keystore");
            System.setProperty("javax.net.ssl.trustStorePassword", "123456");
            ConnectionString connString = new ConnectionString("mongodb://
<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin&replicaSet=replica&ssl=true");
            MongoClientSettings settings = MongoClientSettings.builder()
                    .applyConnectionString(connString)
                    .applyToSslSettings(builder -> builder.enabled(true))
                    .applyToSslSettings(builder -> builder.invalidHostNameAllowed(true))
                    .build();
            MongoClient mongoClient = MongoClients.create(settings);
            MongoDatabase database = mongoClient.getDatabase("admin");
            //Ping the database. If the operation fails, an exception occurs.
            BsonDocument command = new BsonDocument("ping", new BsonInt64(1));
            Document commandResult = database.runCommand(command);
            System.out.println("Connect to database successfully");
        } catch (Exception e) {
            e.printStackTrace();
            System.out.println("Test failed");
        }
    }
}
```

Connection Without the SSL Certificate

NOTE

You do not need to download the SSL certificate because certificate verification on the server is not required.

Connect to a replica set instance using Java. The Java link format is as follows:

```
mongodb://<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin&replicaSet=replica
```

Table 2-28 Parameter description

Parameter	Description
<username>	Current username.
<password>	Password for the current username
<instance_ip>	If you attempt to access the instance from an ECS, set *instance_ip* to the private IP address displayed on the **Basic Information** page of the instance to which you intend to connect.
	If you intend to access the instance through an EIP, set *instance_ip* to the EIP that has been bound to the instance.
<instance_port>	Database port displayed on the **Basic Information** page. Default value: **8635**
<database_name>	Name of the database to be connected.
authSource	Authentication user database. The value is **admin**.

For details about the Java code, see the following example:

```
public class Connector {
    public static void main(String[] args) {
        try {
            ConnectionString connString = new ConnectionString("mongodb://
<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin&replicaSet=replica");
            MongoClientSettings settings = MongoClientSettings.builder()
                    .applyConnectionString(connString)
                    .retryWrites(true)
                    .build();
            MongoClient mongoClient = MongoClients.create(settings);
            MongoDatabase database = mongoClient.getDatabase("admin");
            //Ping the database. If the operation fails, an exception occurs.
            BsonDocument command = new BsonDocument("ping", new BsonInt64(1));
            Document commandResult = database.runCommand(command);
            System.out.println("Connect to database successfully");
        } catch (Exception e) {
            e.printStackTrace();
            System.out.println("Test failed");
        }
    }
}
```

2.2.5.2 Python

This section describes how to connect to a replica set instance using Python.

Prerequisites

1. To connect an ECS to an instance, the ECS must be able to communicate with the DDS instance. You can run the following command to connect to the IP address and port of the instance server to test the network connectivity.

   curl *ip:port*

   If the message **It looks like you are trying to access MongoDB over HTTP on the native driver port** is displayed, the network connectivity is normal.
2. Install Python and third-party installation package **pymongo** on the ECS. Pymongo 2.8 is recommended.
3. If SSL is enabled, you need to download the root certificate and upload it to the ECS.

Connection Code

- Enabling SSL

```
import ssl
from pymongo import MongoClient
conn_urls="mongodb://rwuser:rwuserpassword@ip:port/{mydb}?authSource=admin&replicaSet=replica"
connection = MongoClient(conn_urls,connectTimeoutMS=5000,ssl=True,
ssl_cert_reqs=ssl.CERT_REQUIRED,ssl_match_hostname=False,ssl_ca_certs=${path to certificate authority file})
dbs = connection.database_names()
print "connect database success! database names is %s" % dbs
```

- Disabling SSL

```
import ssl
from pymongo import MongoClient
conn_urls="mongodb://rwuser:rwuserpassword@ip:port/{mydb}?authSource=admin&replicaSet=replica"
connection = MongoClient(conn_urls,connectTimeoutMS=5000)
dbs = connection.database_names()
print "connect database success! database names is %s" % dbs
```

NOTE

- The authentication database in the URL must be **admin**. That means setting **authSource** to **admin**.
- In SSL mode, you need to manually generate the trustStore file.
- The authentication database must be **admin**, and then switch to the service database.

2.2.5.3 PHP

This section describes how to connect to a replica set instance using PHP.

Prerequisites

1. To connect an ECS to a DDS instance, run the following command to connect to the IP address and port of the instance server to test the network connectivity.

   curl ip:port

   If the message **It looks like you are trying to access MongoDB over HTTP on the native driver port** is displayed, the ECS and DDS instance can communicate with each other.

2. If SSL is enabled, you need to download the root certificate and upload it to the ECS.

Obtaining and Using PHP

For the information about PHP, visit **https://www.php.net/mongodb-driver-manager.construct**

Connection Code

- Enabling SSL
  - Run **MongoDB\Client::__construct()** to create a client instance.

```
function __construct(
    ?string $uri = null,
    array $uriOptions = [],
    array $driverOptions = []
)
```

  - Use $uriOptions to set **SSL** to **true** to enable the SSL connection. Use $driverOptions to set **ca_file** to the CA certificate path and **allow_invalid_hostname** to **true**.

```
<?php

require 'vendor/autoload.php'; // include Composer goodies

$replicaset_url = 'mongodb://rwuser:*****@192.168.***.***:8635,192.168.***.***:8635/
test?authSource=admin&replicaSet=replica';
$test_db = 'test_db';
$test_coll = 'test_coll';

//Create mongoclient.
$client = new MongoDB\Client(
....$replicaset_url,
    [
        'ssl' => true,
    ],
    [
        "ca_file" => "/path/to/ca.pem",
        "allow_invalid_hostname" => true
    ]
);

$collection = $client->$test_db->$test_coll;

//Insert a record.
$result = $collection->insertOne([
    'username' => 'admin',
    'email' => 'admin@example.com',
]);
```

```
echo "Object ID: '{$result->getInsertedId()}'", "\n";

//Query a record.
$result = $collection->find(['username' => 'admin']);
foreach ($result as $entry) {
    echo $entry->_id, ': ', $entry->email, "\n";
}

?>
```

- Disabling SSL

```
<?php

require 'vendor/autoload.php'; // include Composer goodies

$replicaset_url = 'mongodb://rwuser:*****@192.168.***.***:8635,192.168.***.***:8635/test?authSource=admin&replicaSet=replica';
$test_db = 'test_db';
$test_coll = 'test_coll';

//Create mongoclient.
$client = new MongoDB\Client($replicaset_url);
$collection = $client->$test_db->$test_coll;

//Insert a record.
$result = $collection->insertOne([
    'username' => 'admin',
    'email' => 'admin@example.com',
]);
echo "Object ID: '{$result->getInsertedId()}'", "\n";

//Query a record.
$result = $collection->find(['username' => 'admin']);
foreach ($result as $entry) {
    echo $entry->_id, ': ', $entry->email, "\n";
}

?>
```

NOTE

- The authentication database in the URL must be **admin**. Set **authSource** to **admin**.
- Change the authentication database of the **rwuser** user to **admin**, and then switch to the service database after authentication.

2.3 Connecting to a Single Node Instance

2.3.1 Connection Methods

You can access DDS over private or public networks.

Table 2-29 Connection methods

Metho d	IP Address	Scenario	Description
DAS	Not required	DAS provides a GUI and allows you to perform visualized operations on the console. SQL execution, advanced database management, and intelligent O&M are available to make database management simple, secure, and intelligent.	• Easy to use, secure, advanced, and intelligent • Recommended
Private netwo rk	Private IP address	DDS provides a private IP address by default. If your applications are running on an ECS in the same region, AZ, and VPC subnet as your DDS instance, you are advised to use a private IP address to connect the ECS to your DDS instances.	Secure and excellent performance
Public netwo rk	EIP	• If your applications are running on an ECS that is in a different region from the one where the DB instance is located, use an EIP to connect the ECS to your DDS DB instances. • If your applications are deployed on another cloud platform, EIP is recommended.	• Low security • For faster transmission and improved security, you are advised to migrate your applications to an ECS that is in the same subnet as your DDS instance and use a private IP address to access the instance.

2.3.2 (Recommended) Connecting to a Single Node Instance Through DAS

2.3.2.1 Overview

DAS provides a GUI and allows you to perform visualized operations on the console. SQL execution, advanced database management, and intelligent O&M are available to make database management simple, secure, and intelligent. You are advised to use DAS to connect to DB instances.

This section describes how to connect to a single node instance through DAS.

Process

To connect to a single node instance, perform the following steps:

1. **Connect to a single node instance through DAS.**

2.3.2.2 Connecting to a Single Node Instance Through DAS

Data Admin Service (DAS) enables you to manage DB instances on a web-based console, simplifying database management and improving working efficiency. You can connect and manage instances through DAS. By default, you have the permission required for remote login. It is recommended that you use the DAS service to connect to instances. DAS is secure and convenient.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

If you want compute and network resources dedicated to your exclusive use, **enable a DeC** and **apply for DCC resources**. After enabling a DeC, you can select the DeC region and project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click **Log In** in the **Operation** column.

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

Figure 2-85 Instance management

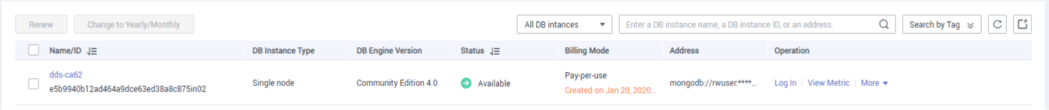

Name/ID	DB Instance Type	DB Engine Version	Status	Billing Mode	Address	Operation
dds-ca62 e5b9940b12ad464a9dce63ed38a8c875in02	Single node	Community Edition 4.0	Available	Pay-per-use Created on Jan 20, 2020...	mongodb://rwuser:****...	Log In | View Metric | More

Step 5 On the displayed login page, enter the administrator username and password and click **Login**.

For details about how to manage databases through DAS, see **Database Management**.

----End

2.3.3 Connecting to a Single Node Instance over a Private Network

2.3.3.1 Configuring a Security Group

A security group is a logical group. It provides access control policies for the ECSs and instances that have the same security protection requirements and are mutually trusted in a VPC.

To ensure database security and reliability, you need to configure security group rules to allow specific IP addresses and ports to access DDS instances.

You can connect to an instance by configuring security group rules in following two ways:

- If the ECS and instance are in the same security group, they can communicate with each other by default. No security group rule needs to be configured. Go to **Connecting to a Single Node Instance Using Mongo Shell (Private Network)**.

Figure 2-86 Same security group

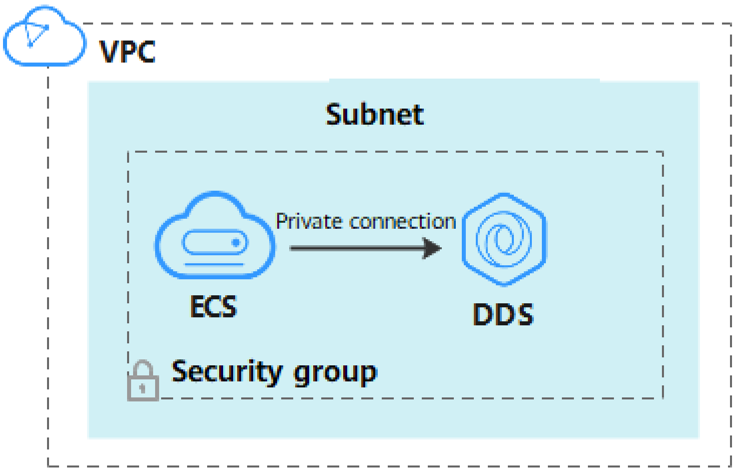

  – If the ECS and instance are in different security groups, you need to configure security group rules for them, separately.

Figure 2-87 Different security groups

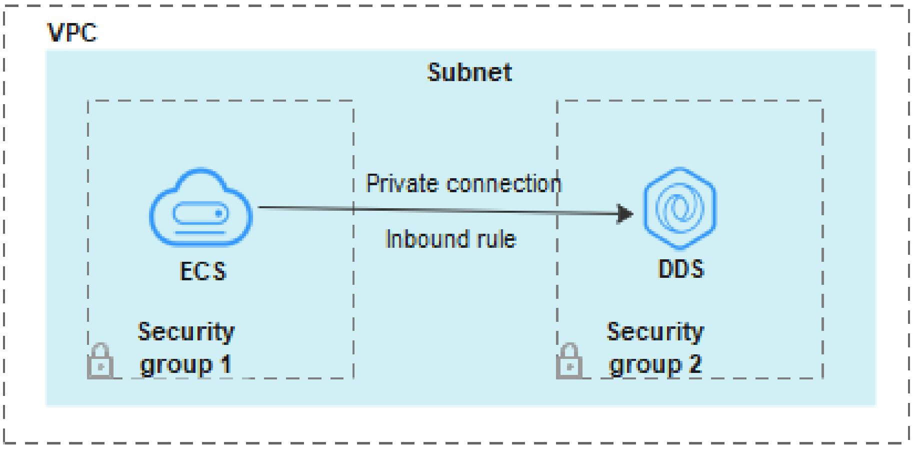

- Instance: Configure an **inbound rule** for the security group associated with the instance.
- ECS: The default security group rule allows all outbound data packets. In this case, you do not need to configure a security group rule for the ECS. If not all traffic is allowed to reach the instance, configure an **outbound** rule for the ECS.

This section describes how to configure an inbound rule for an instance.

Precautions

- By default, an account can create up to 500 security group rules.
- Too many security group rules will increase the first packet latency, so a maximum of 50 rules for each security group is recommended.
- By default, one DDS instance is associated with only one security group.
- DDS allows you to associate multiple security groups to a DB instance. You can apply for the service based on your service requirements. For better network performance, you are advised to select no more than five security groups.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Network Information** area on the **Basic Information** page, click the security group.

Figure 2-88 Security Group

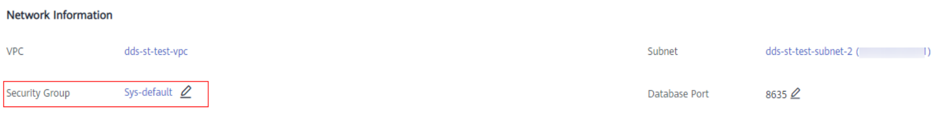

Network Information

VPC	dds-st-test-vpc	Subnet	dds-st-test-subnet-2 ([redacted])
Security Group	Sys-default	Database Port	8635

You can also choose **Connections** in the navigation pane on the left. On the **Private Connection** tab, in the **Security Group** area, click the security group name.

Figure 2-89 Security Group

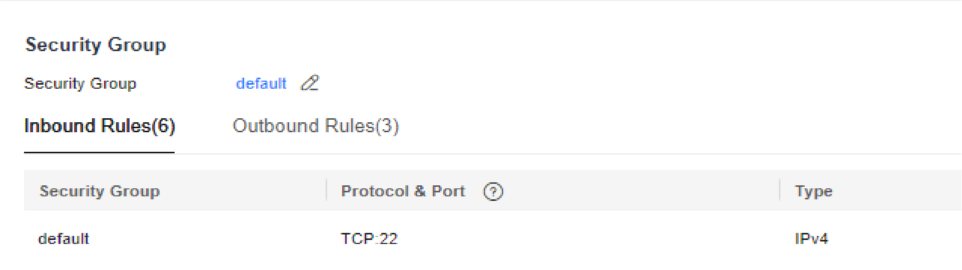

Security Group

Security Group default

Inbound Rules(6) Outbound Rules(3)

Security Group	Protocol & Port	Type
default	TCP:22	IPv4

Step 6 On the **Security Group** page, locate the target security group and click **Manage Rule** in the **Operation** column.

Step 7 On the **Inbound Rules** tab, click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Step 8 Add a security group rule as prompted.

Figure 2-90 Add Inbound Rule

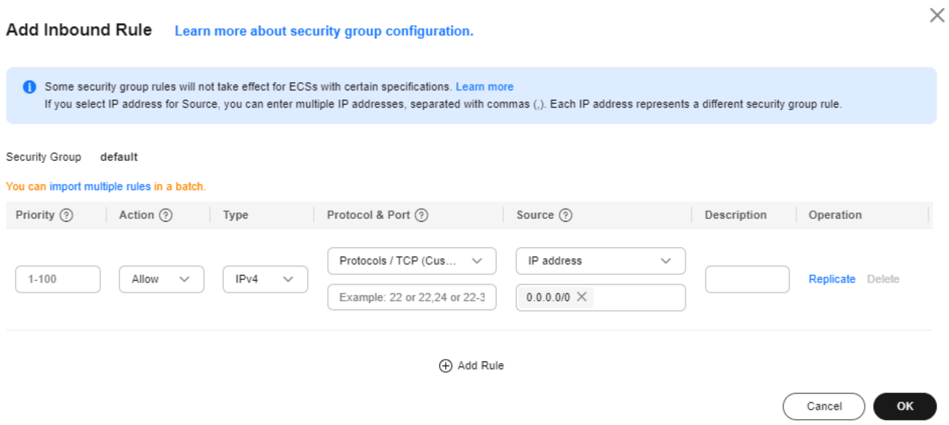

Add Inbound Rule Learn more about security group configuration.

Some security group rules will not take effect for ECSs with certain specifications. Learn more
If you select IP address for Source, you can enter multiple IP addresses, separated with commas (,). Each IP address represents a different security group rule.

Security Group default

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Source	Description	Operation
1-100	Allow	IPv4	Protocols / TCP (Cus... Example: 22 or 22,24 or 22-3	IP address 0.0.0.0/0		Replicate Delete

Table 2-30 Inbound rule settings

Parameter	Description	Example
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default priority is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1
Action	The security group rule actions. A rule with a deny action overrides another with an allow action if the two rules have the same priority.	Allow

Parameter	Description	Example
Protocol & Port	The network protocol required for access. Available options: **TCP**, **UDP**, **ICMP**, or **GRE**	TCP
	Port: the port on which you wish to allow access to DDS. The default port is 8635. The port ranges from 2100 to 9500 or can be 27017, 27018, or 27019.	8635
Type	IP address type. Only **IPv4** and **IPv6** are supported.	IPv4
Source	Specifies the supported IP address, security group, and IP address group, which allow access from IP addresses or instances in other security group. Example: ● Single IP address: 192.168.10.10/32 ● IP address segment: 192.168.1.0/24 ● All IP addresses: 0.0.0.0/0 ● Security group: sg-abc ● IP address group: ipGroup-test If you enter a security group, all ECSs associated with the security group comply with the created rule. For more information about IP address groups, see **IP Address Group Overview**.	0.0.0.0/0
Description	(Optional) Provides supplementary information about the security group rule. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

Step 9 Click **OK**.

----End

2.3.3.2 Connecting to a Single Node Instance Using Mongo Shell (Private Network)

Mongo shell is the default client for the MongoDB database server. You can use Mongo Shell to connect to DB instances, and query, update, and manage data in databases. DDS is compatible with MongoDB. Mongo Shell is a part of the MongoDB client. To use Mongo Shell, download and install the MongoDB client first, and then use the Mongo shell to connect to the DB instance.

By default, a DDS instance provides a private IP address. If your applications are deployed on an ECS and are in the same region and VPC as DDS instances, you can connect to DDS instances using a private IP address to achieve a fast transmission rate and high security.

This section describes how to use Mongo Shell installed on a Linux ECS to connect to a single node instance over a private network.

You can connect to an instance using an SSL connection or an unencrypted connection. The SSL connection is encrypted and more secure. To improve data transmission security, connect to instances using SSL.

Prerequisites

1. For details about how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS**.
2. Install the MongoDB client on the ECS.

   For details about how to install a MongoDB client, see **How Can I Install a MongoDB Client?**
3. The ECS can communicate with the DDS instance. For details, see **Configuring a Security Group**.

SSL

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Connections**.

Step 6 In the **Basic Information** area, click 📥 next to the **SSL** field.

Step 7 Import the root certificate to the Linux or Windows ECS. For details, see **How Can I Import the Root Certificate to a Windows or Linux OS?**

Step 8 Connect to a DDS instance.

Using a private IP address

Example command:

```
./mongo --host <DB_HOST> --port <DB_PORT> -u <DB_USER> -p --authenticationDatabase admin --ssl --sslCAFile<FILE_PATH> --sslAllowInvalidHostnames
```

Parameter description:

- **DB_HOST** is the private IP address of the instance to be connected.

  On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. On the **Private Connection** tab, obtain the IP address of the corresponding node.

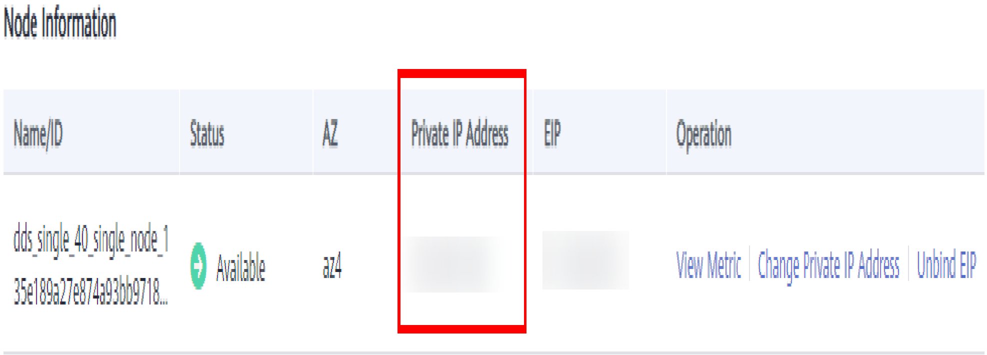

Node Information

Name/ID	Status	AZ	Private IP Address	EIP	Operation
dds_single_40_single_node_1 35e189a27e874a93bb9718...	Available	az4	[redacted]	[redacted]	View Metric | Change Private IP Address | Unbind EIP

- **DB_PORT** is the database port. The default port number is 8635.

  You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Private Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

  Figure 2-91 Obtaining the port

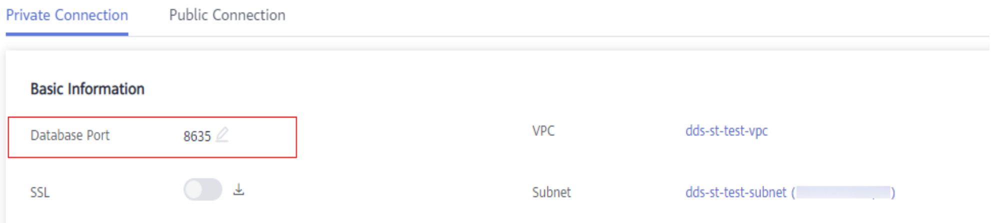

Basic Information

Database Port	8635	VPC	dds-st-test-vpc
SSL		Subnet	dds-st-test-subnet ([redacted])

- **DB_USER** is the database user. The default value is **rwuser**.
- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: To ensure that the internal communication of the single node instance does not occupy resources such as the user IP address and bandwidth, the single node certificate is generated using the internal management IP address. **--sslAllowInvalidHostnames** is needed for the SSL connection through a private network.

Command example:

```
./mongo --host 192.168.xx.xx --port 8635 -u rwuser -p --authenticationDatabase admin --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames
```

Enter the database account password when prompted:

```
Enter password:
```

Step 9 Check the connection result. If the following information is displayed, the connection is successful.

```
replica:PRIMARY>
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Log in to the ECS.

Step 2 Connect to a DDS instance.

Using a private IP address

Example command:

./mongo --host<*DB_HOST*>**--port**<*DB_PORT*>**-u**<*DB_USER*>**-p --authenticationDatabase admin**

Parameter description:

- **DB_HOST** is the private IP address of the instance to be connected.

  On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections**. On the **Private Connection** tab, obtain the IP address of the corresponding node.

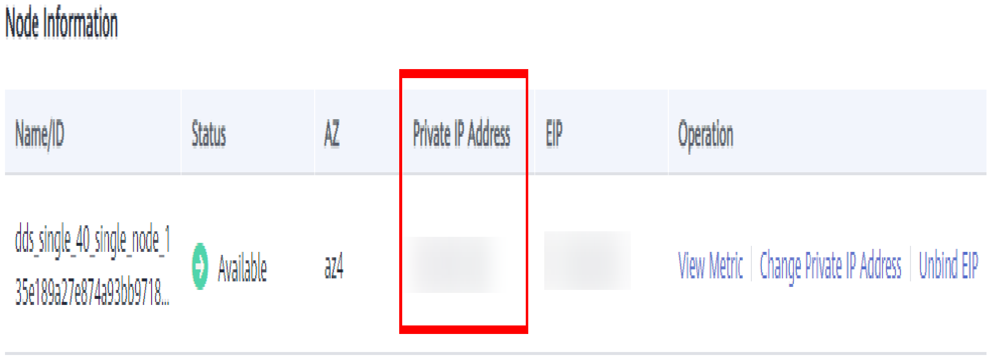

Node Information

Name/ID	Status	AZ	Private IP Address	EIP	Operation
dds_single_40_single_node_1 35e189a27e874a93bb9718...	Available	az4	[redacted]	[redacted]	View Metric Change Private IP Address Unbind EIP

- **DB_PORT** is the database port. The default port number is 8635.

  You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Private Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

Figure 2-92 Obtaining the port

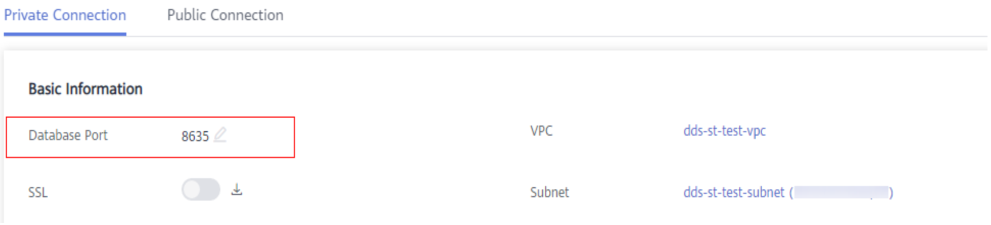

- **DB_USER** is the database user. The default value is **rwuser**.

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin**

Enter the database account password when prompted:

```
Enter password:
```

Step 3 Check the connection result. If the following information is displayed, the connection is successful.

```
replica:PRIMARY>
```

----End

2.3.4 Connecting to a Single Node Instance over a Public Network

2.3.4.1 Binding an EIP

After you create an instance, you can bind an EIP to it to allow external access. If later you want to prohibit external access, you can also unbind the EIP from the instance.

Precautions

- Deleting a bound EIP does not mean that the EIP is unbound.
- Before accessing a database, apply for an EIP on the VPC console. Then, add an inbound rule to allow the IP addresses or IP address ranges of ECSs. For details, see **Configuring a Security Group**.
- To change the EIP that has been bound to a node, unbind it from the node first.

Binding an EIP

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

Step 2 Click ⊙ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the single node instance name.

Step 5 In the navigation pane on the left, choose **Connections**. Click the **Public Connection** tab. In the **Basic Information** area, locate the node you want to bind an EIP to and click **Bind EIP** in the **Operation** column.

Figure 2-93 Binding an EIP

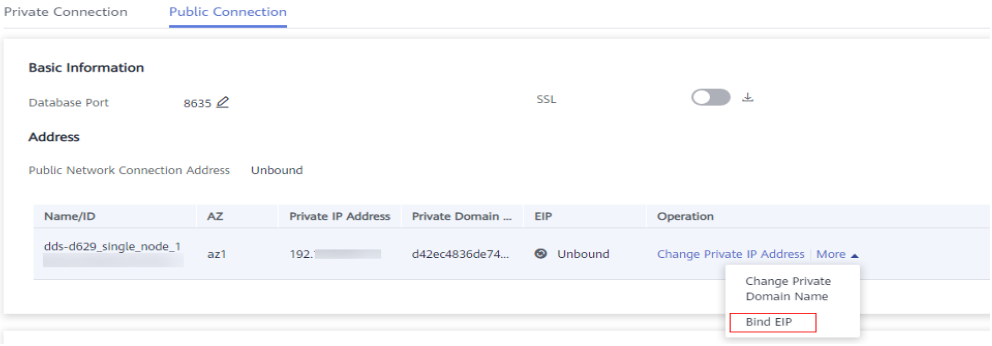

Private Connection | Public Connection

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address Unbound

Name/ID	AZ	Private IP Address	Private Domain ...	EIP	Operation
dds-d629_single_node_1 [redacted]	az1	192.[redacted]	d42ec4836de74...	Unbound	Change Private IP Address | More

You can also locate the node in the **Node Information** area on the **Basic Information** page and click **Bind EIP** in the **Operation** column.

Figure 2-94 Binding an EIP

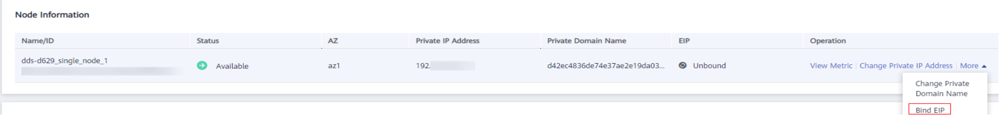

Node Information

Name/ID	Status	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-d629_single_node_1 [redacted]	Available	az1	192.[redacted]	d42ec4836de74e37ae2e19da03...	Unbound	View Metric | Change Private IP Address | More

Step 6 In the displayed dialog box, all available unbound EIPs are listed. Select the required EIP and click **OK**. If no available EIPs are displayed, click **View EIP** and create an EIP on the VPC console.

Figure 2-95 Selecting an EIP

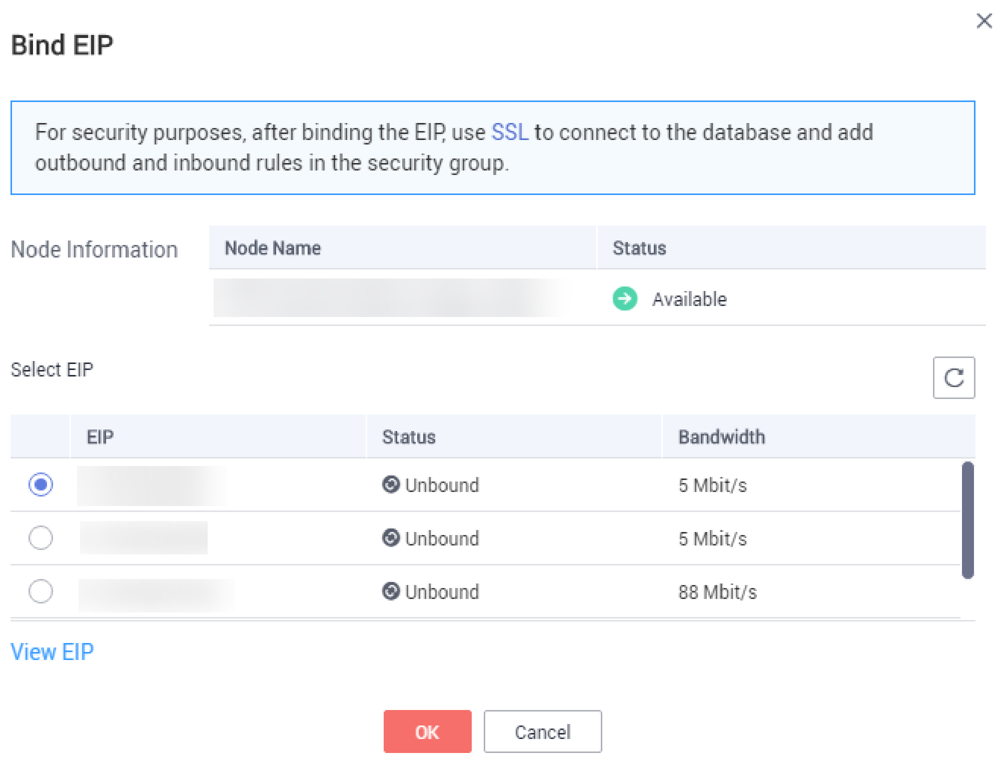

Bind EIP

For security purposes, after binding the EIP, use SSL to connect to the database and add outbound and inbound rules in the security group.

Node Information

Node Name	Status
[redacted]	Available

Select EIP

	EIP	Status	Bandwidth
●	[redacted]	Unbound	5 Mbit/s
○	[redacted]	Unbound	5 Mbit/s
○	[redacted]	Unbound	88 Mbit/s

Step 7 In the **EIP** column, you can view the EIP that was bound.

To unbind an EIP from the instance, see **Unbinding an EIP**.

----End

Unbinding an EIP

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the single node instance name.

Step 5 In the navigation pane on the left, choose **Connections**. Click the **Public Connection** tab. In the **Basic Information** area, locate the node and click **Unbind EIP** in the **Operation** column.

Figure 2-96 Unbinding an EIP

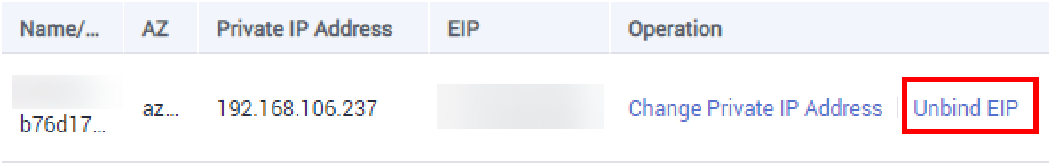

Name/...	AZ	Private IP Address	EIP	Operation
[redacted] b76d17...	az...	192.168.106.237	[redacted]	Change Private IP Address Unbind EIP

You can also locate the node in the **Node Information area** on the **Basic Information** page and click **Unbind EIP** in the **Operation** column.

Step 6 In the displayed dialog box, click **Yes**.

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

----End

2.3.4.2 Configuring a Security Group

A security group is a logical group. It provides access control policies for the ECSs and instances that have the same security protection requirements and are mutually trusted in a VPC.

To ensure database security and reliability, you need to configure security group rules to allow specific IP addresses and ports to access DDS instances.

If you attempt to connect to an instance through an EIP, you need to configure an inbound rule for the security group associated with the instance.

Precautions

- By default, an account can create up to 500 security group rules.
- Too many security group rules will increase the first packet latency, so a maximum of 50 rules for each security group is recommended.
- By default, one DDS instance is associated with only one security group.
- DDS allows you to associate multiple security groups to a DB instance. You can apply for the service based on your service requirements. For better network performance, you are advised to select no more than five security groups.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Network Information** area on the **Basic Information** page, click the security group.

Figure 2-97 Security Group

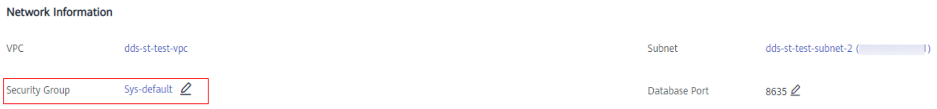

Network Information

VPC dds-st-test-vpc

Subnet dds-st-test-subnet-2 ([redacted])

Security Group Sys-default

Database Port 8635

You can also choose **Connections** in the navigation pane on the left. On the **Private Connection** tab, in the **Security Group** area, click the security group name.

Figure 2-98 Security Group

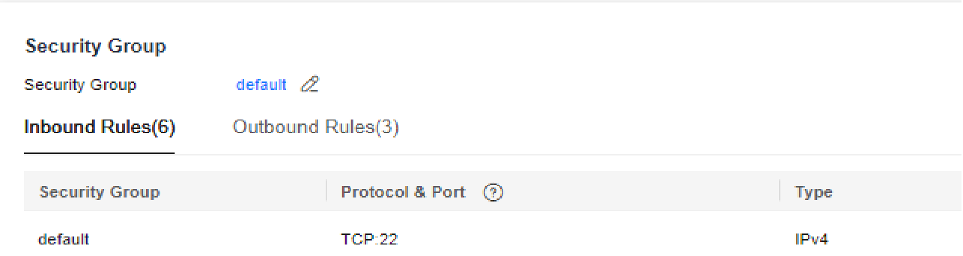

Security Group

Security Group default

Inbound Rules(6) Outbound Rules(3)

Security Group	Protocol & Port	Type
default	TCP:22	IPv4

Step 6 On the **Security Group** page, locate the target security group and click **Manage Rule** in the **Operation** column.

Step 7 On the **Inbound Rules** tab, click **Add Rule**. The **Add Inbound Rule** dialog box is displayed.

Step 8 Add a security group rule as prompted.

Figure 2-99 Add Inbound Rule

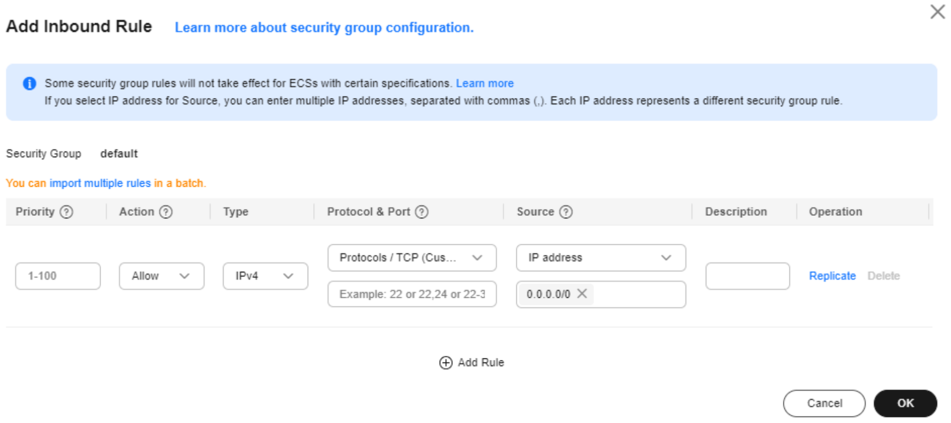

Add Inbound Rule Learn more about security group configuration.

Some security group rules will not take effect for ECSs with certain specifications. Learn more
If you select IP address for Source, you can enter multiple IP addresses, separated with commas (,). Each IP address represents a different security group rule.

Security Group default

You can import multiple rules in a batch.

Priority	Action	Type	Protocol & Port	Source	Description	Operation
1-100	Allow	IPv4	Protocols / TCP (Cus... Example: 22 or 22,24 or 22-3	IP address 0.0.0.0/0		Replicate Delete

Table 2-31 Inbound rule settings

Parameter	Description	Example
Priority	The security group rule priority. The priority value ranges from 1 to 100. The default priority is 1 and has the highest priority. The security group rule with a smaller value has a higher priority.	1

Parameter	Description	Example
Action	The security group rule actions. A rule with a deny action overrides another with an allow action if the two rules have the same priority.	Allow
Protocol & Port	The network protocol required for access. Available options: **TCP**, **UDP**, **ICMP**, or **GRE**	TCP
	Port: the port on which you wish to allow access to DDS. The default port is 8635. The port ranges from 2100 to 9500 or can be 27017, 27018, or 27019.	8635
Type	IP address type. Only **IPv4** and **IPv6** are supported.	IPv4
Source	Specifies the supported IP address, security group, and IP address group, which allow access from IP addresses or instances in other security group. Example: ● Single IP address: 192.168.10.10/32 ● IP address segment: 192.168.1.0/24 ● All IP addresses: 0.0.0.0/0 ● Security group: sg-abc ● IP address group: ipGroup-test If you enter a security group, all ECSs associated with the security group comply with the created rule. For more information about IP address groups, see **IP Address Group Overview**.	0.0.0.0/0
Description	(Optional) Provides supplementary information about the security group rule. This parameter is optional. The description can contain a maximum of 255 characters and cannot contain angle brackets (< or >).	-

Step 9 Click **OK**.

----End

2.3.4.3 Connecting to a Single Node Instance Using Mongo Shell (Public Network)

In the following scenarios, you can access a DDS instance from the Internet by binding an EIP to the instance.

Scenario 1: Your applications are running on an ECS that is in a different region from the one where the DDS instance is located.

Figure 2-100 Accessing DDS from ECS across regions

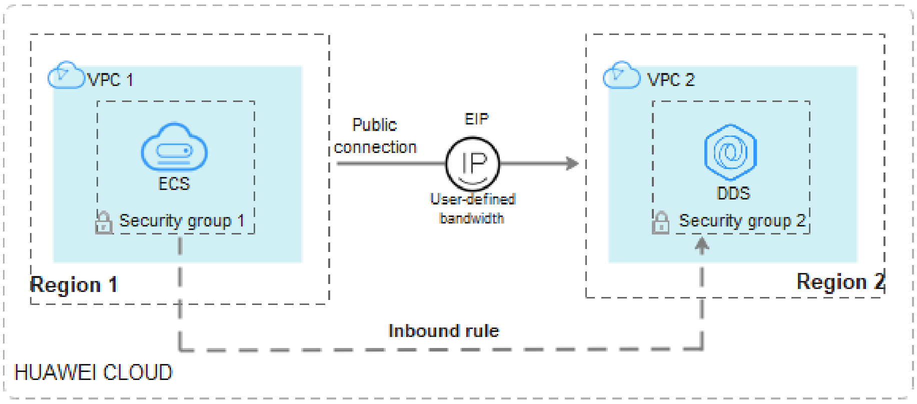

Scenario 2: Your applications are deployed on a cloud server provided by other vendors.

Figure 2-101 Accessing DDS from other cloud servers

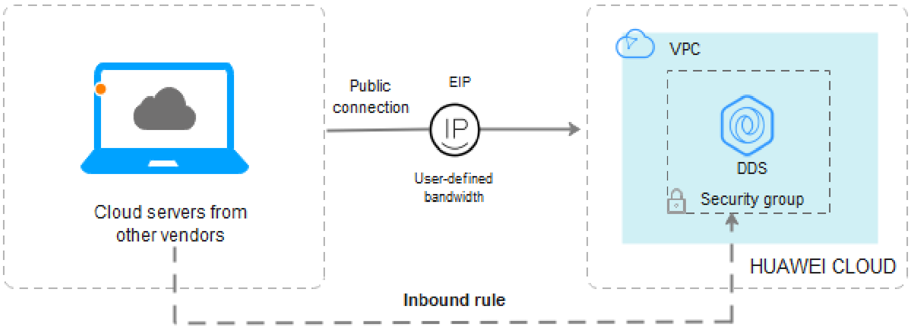

This section describes how to use Mongo Shell to connect to a single node instance through an EIP.

You can connect to an instance using an SSL connection or an unencrypted connection. The SSL connection is encrypted and more secure. To improve data transmission security, connect to instances using SSL.

Prerequisites

1. For details about how to create and log in to an ECS, see **Purchasing an ECS** and **Logging In to an ECS**.
2. **Bind an EIP** to the single node instance and **configure security group rules** to ensure that the EIP can be accessed from the ECS.

3. Install the MongoDB client on the ECS.

   For details about how to install a MongoDB client, see **How Can I Install a MongoDB Client?**

SSL

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Connections**.

Step 6 In the **Basic Information** area, click ⤓ next to the **SSL** field.

Step 7 Import the root certificate to the Linux or Windows ECS. For details, see **How Can I Import the Root Certificate to a Windows or Linux OS?**

Step 8 Connect to the instance in the directory where the MongoDB client is located.

Using an EIP

Example command:

```
./mongo --host <DB_HOST> --port <DB_PORT> -u <DB_USER> -p --authenticationDatabaseadmin --ssl --sslCAFile<FILE_PATH> --sslAllowInvalidHostnames
```

Parameter description:

- **DB_HOST** is the EIP bound to the instance to be connected.

  On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections** > **Public Connection** and obtain the EIP of the corresponding node.

Figure 2-102 Obtaining an EIP

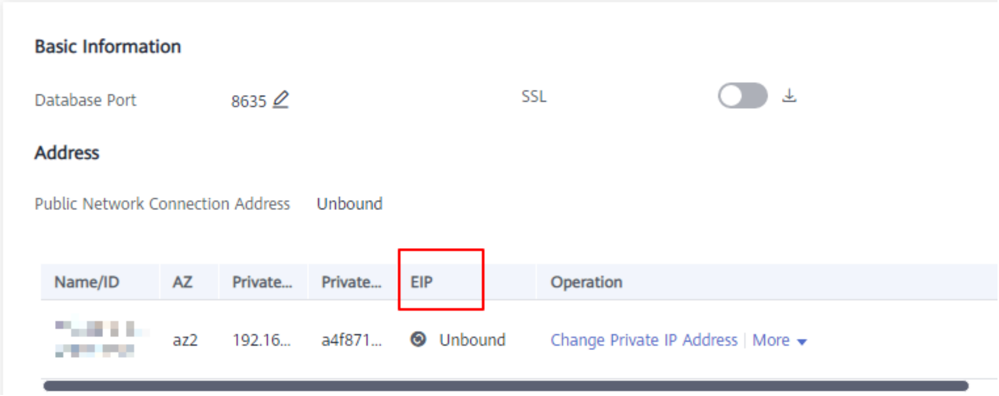

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address Unbound

Name/ID	AZ	Private...	Private...	EIP	Operation
[redacted]	az2	192.16...	a4f871...	Unbound	Change Private IP Address | More

- **DB_PORT** is the database port. The default port number is 8635.

  You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Public Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

Figure 2-103 Obtaining the port

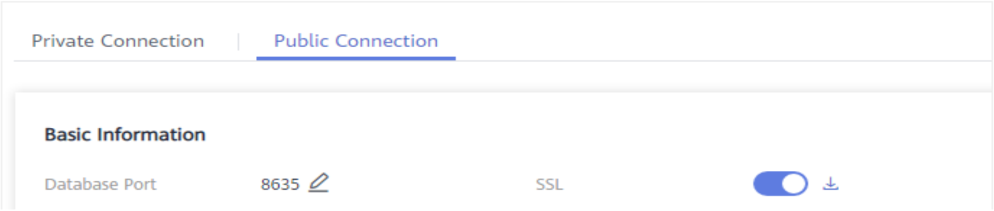

Basic Information

Database Port 8635

SSL

- **DB_USER** is the database user. The default value is **rwuser**.
- **FILE_PATH** is the path for storing the root certificate.
- **--sslAllowInvalidHostnames**: To ensure that the internal communication of the single node instance does not occupy resources such as the user IP address and bandwidth, the single node certificate is generated using the internal management IP address. **--sslAllowInvalidHostnames** is needed for the SSL connection through a public network.

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidHostnames**

Enter the database account password when prompted:

```
Enter password:
```

Step 9 Check the connection result. If the following information is displayed, the connection is successful.

```
replica:PRIMARY>
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Log in to the ECS.

Step 2 Connect to a DDS instance.

Using an EIP

Example command:

./mongo --host *<DB_HOST>* **--port** *<DB_PORT>* **-u** *<DB_USER>* **-p --authenticationDatabase admin**

Parameter description:

- **DB_HOST** is the EIP bound to the instance to be connected.

  On the **Instances** page, click the instance name. The **Basic Information** page is displayed. Choose **Connections** > **Public Connection** and obtain the EIP of the corresponding node.

Figure 2-104 Obtaining an EIP

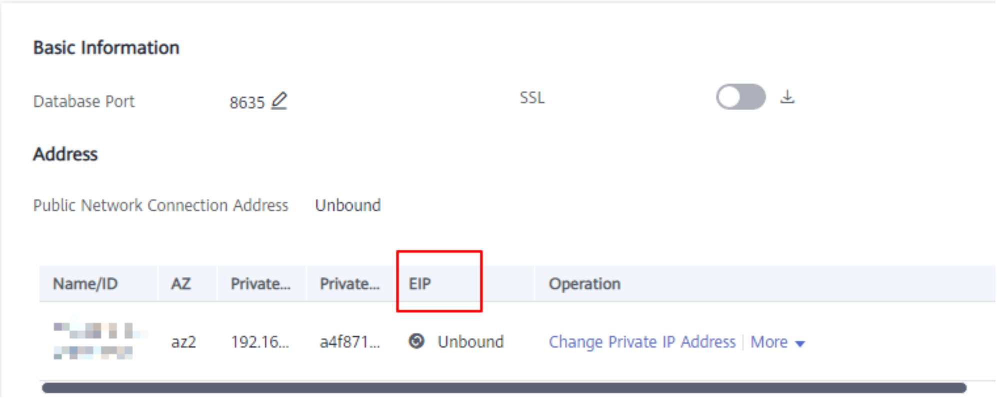

Basic Information

Database Port 8635

SSL

Address

Public Network Connection Address Unbound

Name/ID	AZ	Private...	Private...	EIP	Operation
[redacted]	az2	192.16...	a4f871...	Unbound	Change Private IP Address | More

- **DB_PORT** is the database port. The default port number is 8635.

  You can click the instance name to go to the **Basic Information** page. In the navigation pane on the left, choose **Connections**. On the displayed page, click the **Public Connection** tab and obtain the port from the **Database Port** field in the **Basic Information** area.

Figure 2-105 Obtaining the port

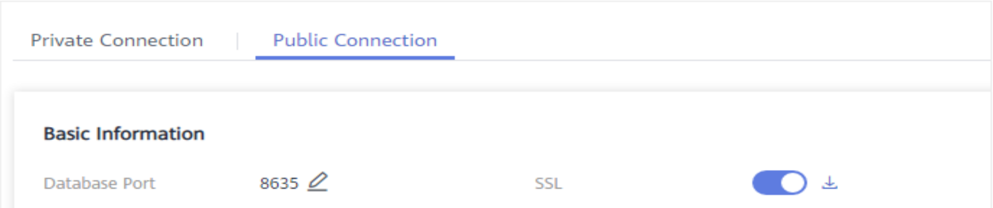

Basic Information

Database Port 8635

SSL

- **DB_USER** is the database user. The default value is **rwuser**.

Command example:

./mongo --host *192.168.xx.xx* **--port 8635 -u rwuser -p --authenticationDatabase admin**

Enter the database account password when prompted:

```
Enter password:
```

Step 3 Check the connection result. If the following information is displayed, the connection is successful.

```
replica:PRIMARY>
```

----End

2.3.4.4 Connecting to a Single Node Instance Using Robo 3T

If you want to connect to an instance from a local device, you can bind an EIP to the instance and use Robo 3T to connect to the instance over a public network.

This section describes how to use Robo 3T to connect to a single node instance from a local device. In this section, the Windows operating system (OS) used by the client is used as an example.

Robo 3T can connect to an instance with an unencrypted connection or an encrypted connection (SSL). To improve data transmission security, connect to instances using SSL.

Connection Diagram

Figure 2-106 Connection diagram

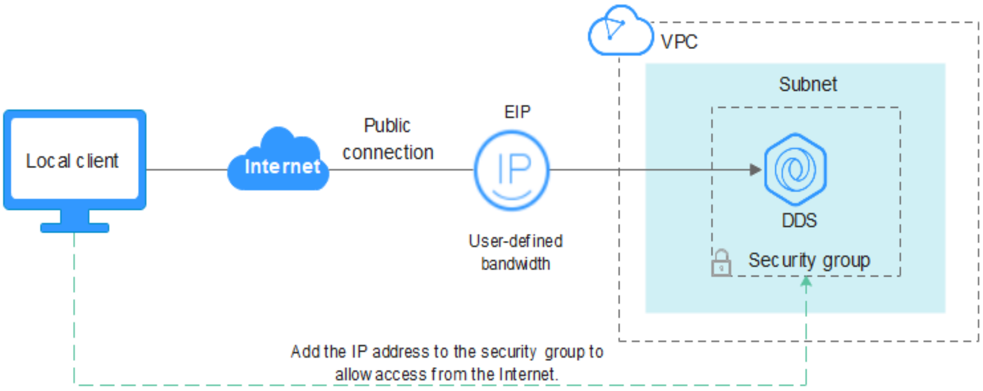

Prerequisites

1. **Bind an EIP** to the single node instance and **configure security group rules** to ensure that the EIP can be accessed using Robo 3T.
2. Install Robo 3T.

   For details about how to install Robo 3T, see **How Can I Install Robo 3T?**

SSL

NOTICE

If you connect to an instance over the SSL connection, enable SSL first. Otherwise, an error is reported. For details about how to enable SSL, see **Enabling and Disabling SSL**.

Step 1 Run the installed Robo 3T. On the displayed dialog box, click **Create**.

Figure 2-107 Connections

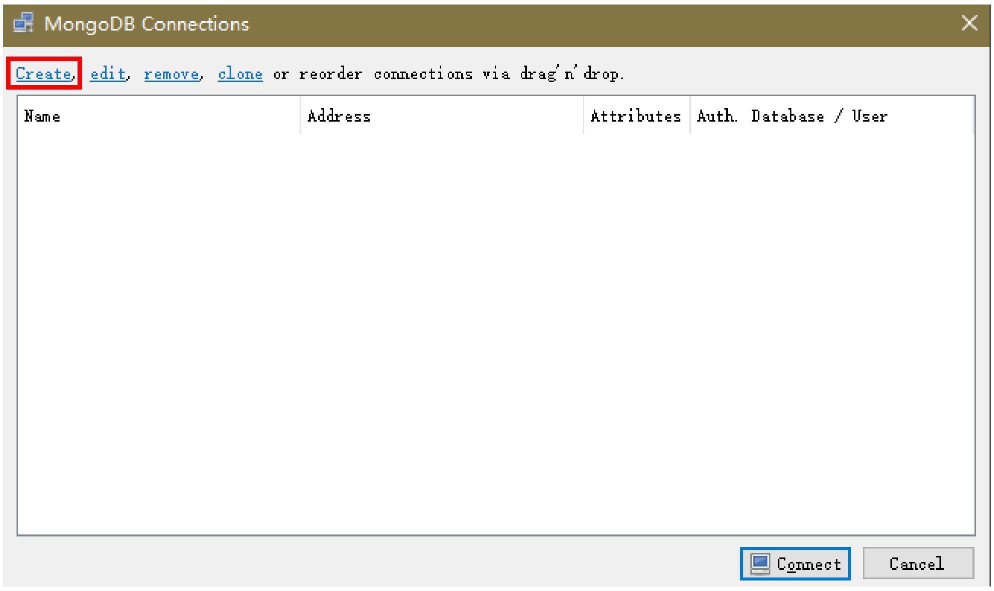

Step 2 In the **Connection Settings** dialog box, set the parameters of the new connection.

1. On the **Connection** tab, enter the name of the new connection in the **Name** text box and enter the EIP and database port that are bound to the DDS DB instance in the **Address** text box.

Figure 2-108 Connection

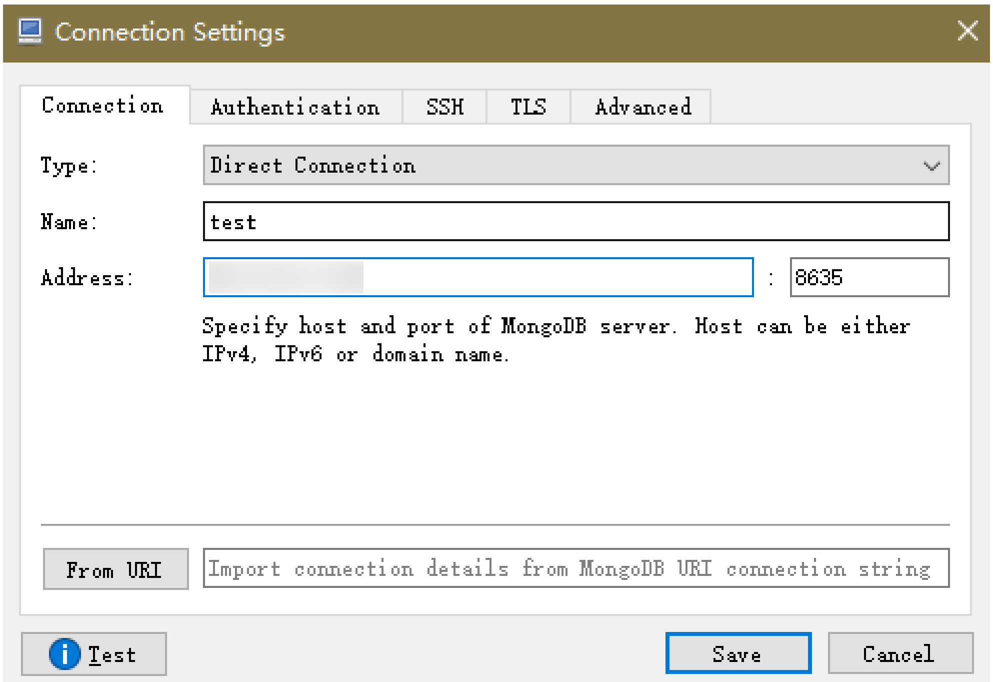

Connection Settings

Type: Direct Connection

Name: test

Address: [redacted] : 8635

Specify host and port of MongoDB server. Host can be either IPv4, IPv6 or domain name.

Import connection details from MongoDB URI connection string

2. On the **Authentication** tab, set **Database** to **admin**, **User Name** to **rwuser**, and **Password** to the administrator password you set during the creation of the cluster instance.

Figure 2-109 Authentication

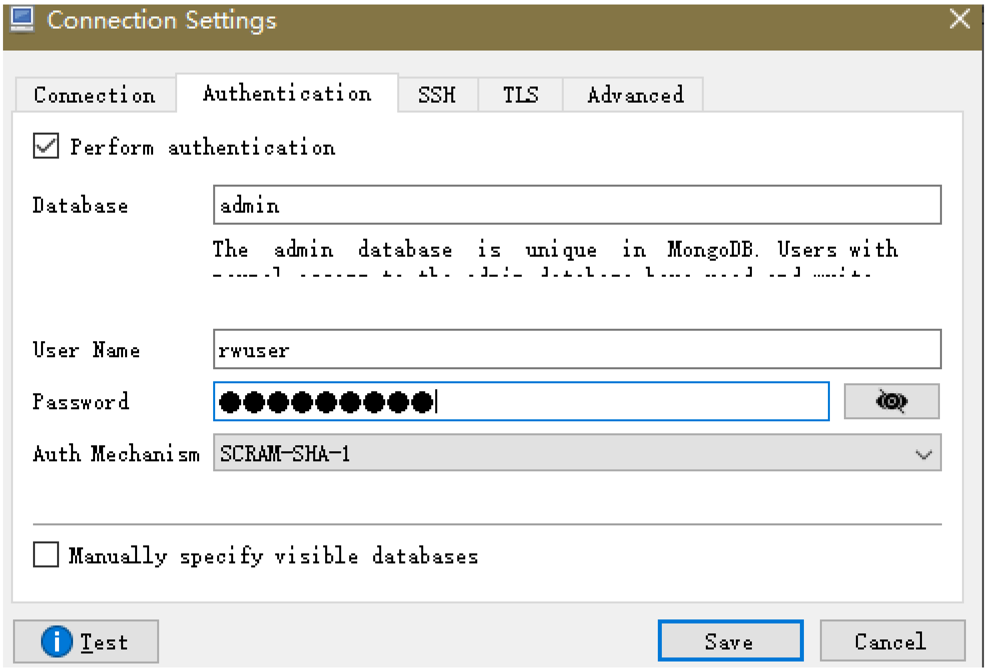

Connection Settings

☑ Perform authentication

Database: admin

The admin database is unique in MongoDB. Users with

User Name: rwuser

Password: ●●●●●●●●●●

Auth Mechanism: SCRAM-SHA-1

☐ Manually specify visible databases

3. On the **TLS** tab, select **Use TLS protocol** and select **Self-signed Certificate** for **Authentication Method**.

Figure 2-110 SSL

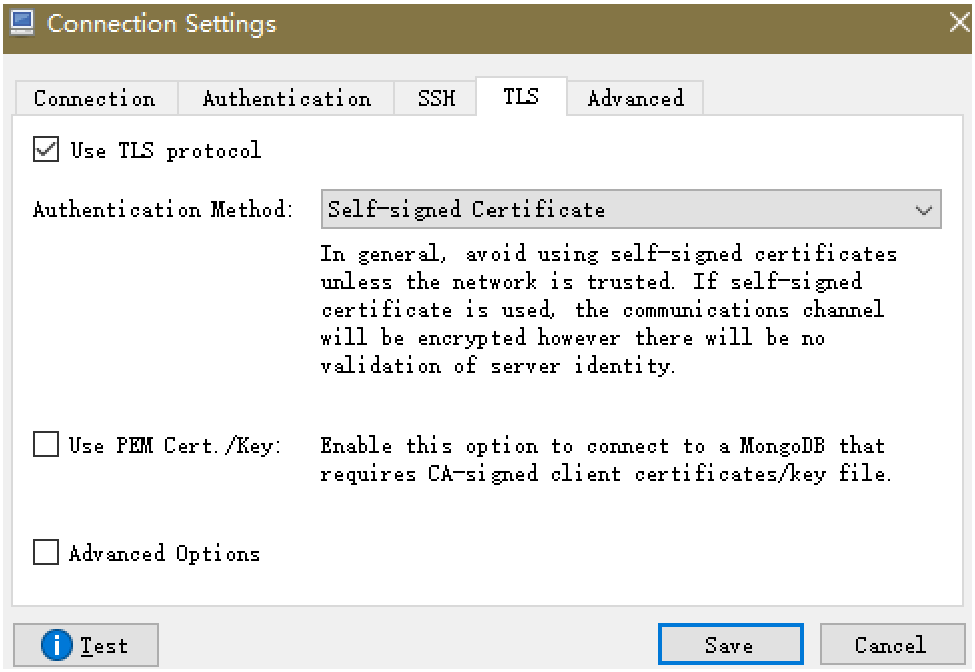

☑ Use TLS protocol

Authentication Method: Self-signed Certificate

In general, avoid using self-signed certificates unless the network is trusted. If self-signed certificate is used, the communications channel will be encrypted however there will be no validation of server identity.

□ Use PEM Cert./Key: Enable this option to connect to a MongoDB that requires CA-signed client certificates/key file.

□ Advanced Options

4. Click **Save**.

Step 3 On the **MongoDB Connections** page, click **Connect** to connect to the single node instance.

Figure 2-111 Single node connection information

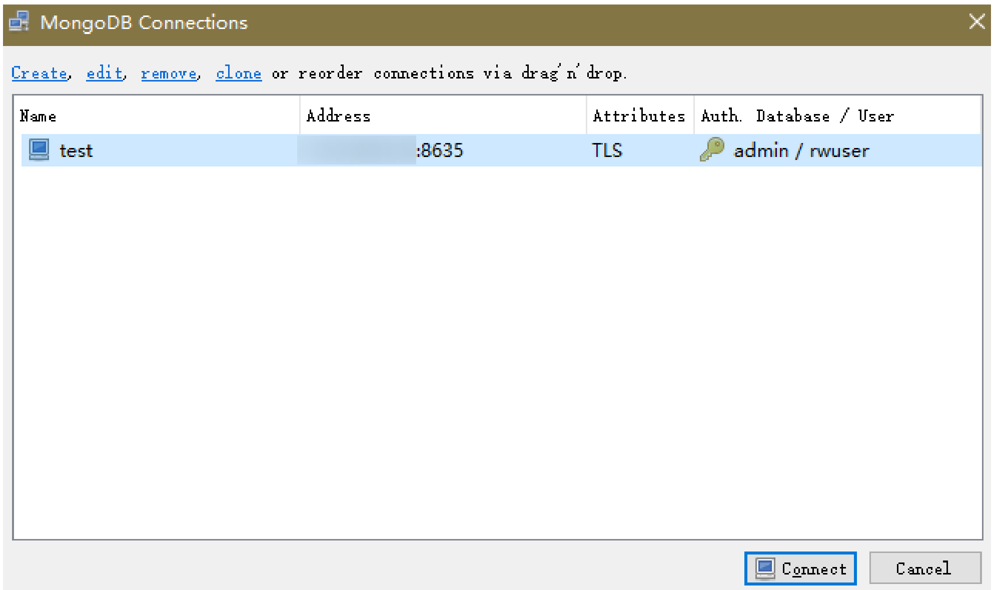

Create, edit, remove, clone or reorder connections via drag'n'drop.

Name	Address	Attributes	Auth. Database / User
test	[redacted]:8635	TLS	admin / rwuser

Step 4 If the single node instance is successfully connected, the page shown in **Figure 2-112** is displayed.

Figure 2-112 Single node connected

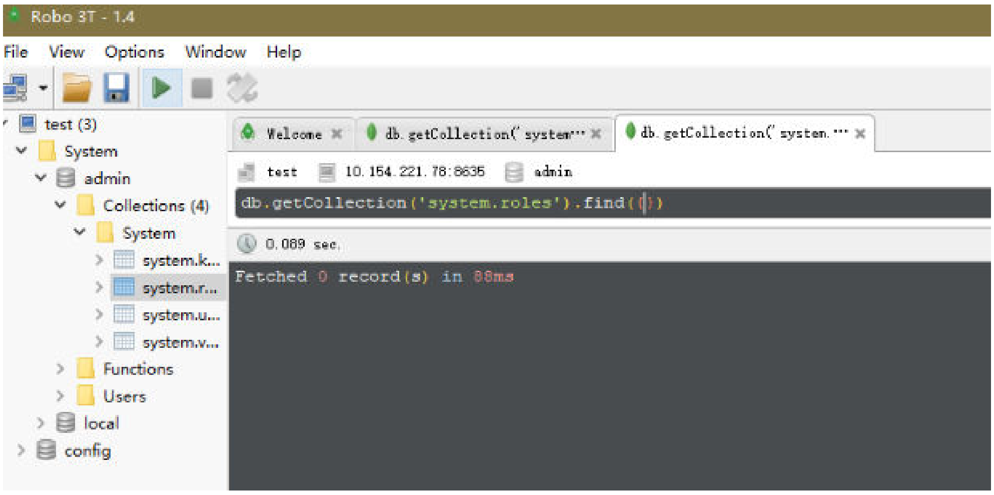

```
db.getCollection('system.roles').find({})
```

0.089 sec.

```
Fetched 0 record(s) in 88ms
```

----End

Unencrypted Connection

NOTICE

If you connect to an instance over an unencrypted connection, disable SSL first. Otherwise, an error is reported. For details about how to disable SSL, see **Enabling and Disabling SSL**.

Step 1 Run the installed Robo 3T. On the displayed dialog box, click **Create**.

Figure 2-113 Connections

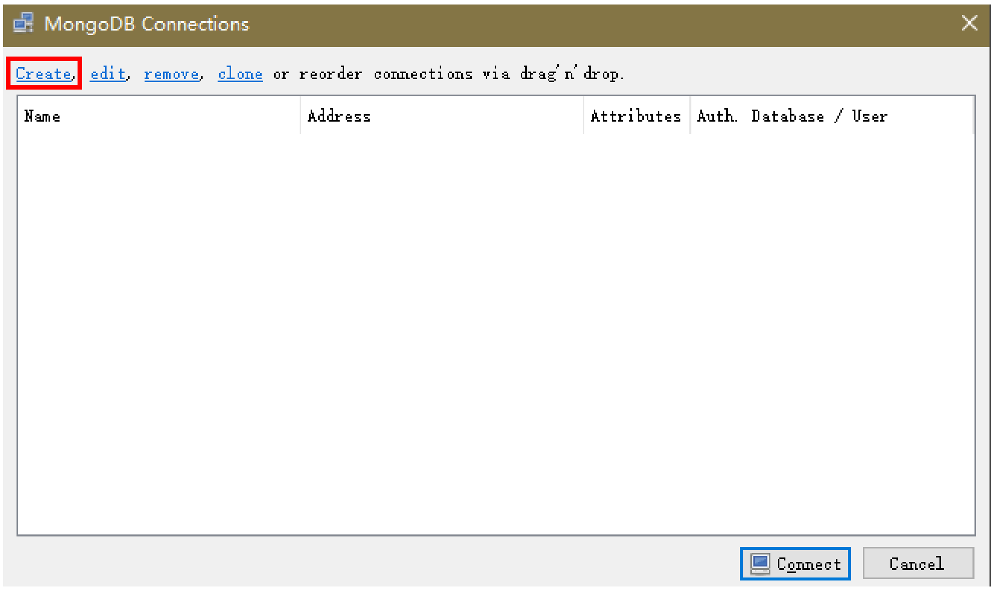

Create, edit, remove, clone or reorder connections via drag'n'drop.

Name	Address	Attributes	Auth. Database / User

Step 2 In the **Connection Settings** dialog box, set the parameters of the new connection.

1. On the **Connection** tab, enter the name of the new connection in the **Name** text box and enter the EIP and database port that are bound to the DDS DB instance in the **Address** text box.

Figure 2-114 Connection

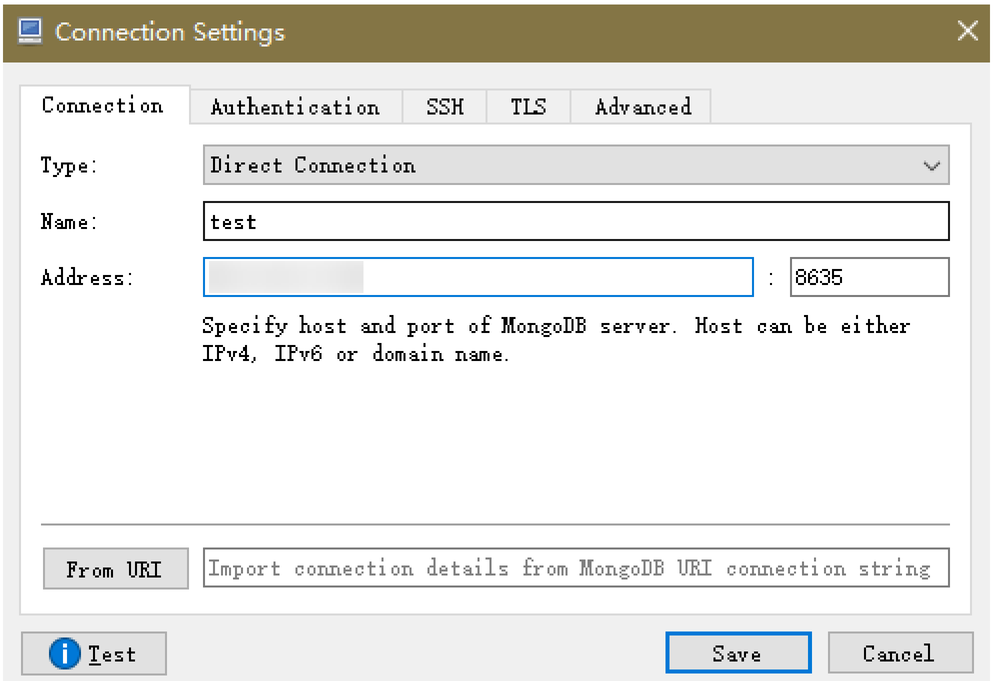

Connection Settings

Connection | Authentication | SSH | TLS | Advanced

Type: Direct Connection

Name: test

Address: [redacted] : 8635

Specify host and port of MongoDB server. Host can be either IPv4, IPv6 or domain name.

From URI: Import connection details from MongoDB URI connection string

2. On the **Authentication** tab, set **Database** to **admin**, **User Name** to **rwuser**, and **Password** to the administrator password you set during the creation of the cluster instance.

Figure 2-115 Authentication

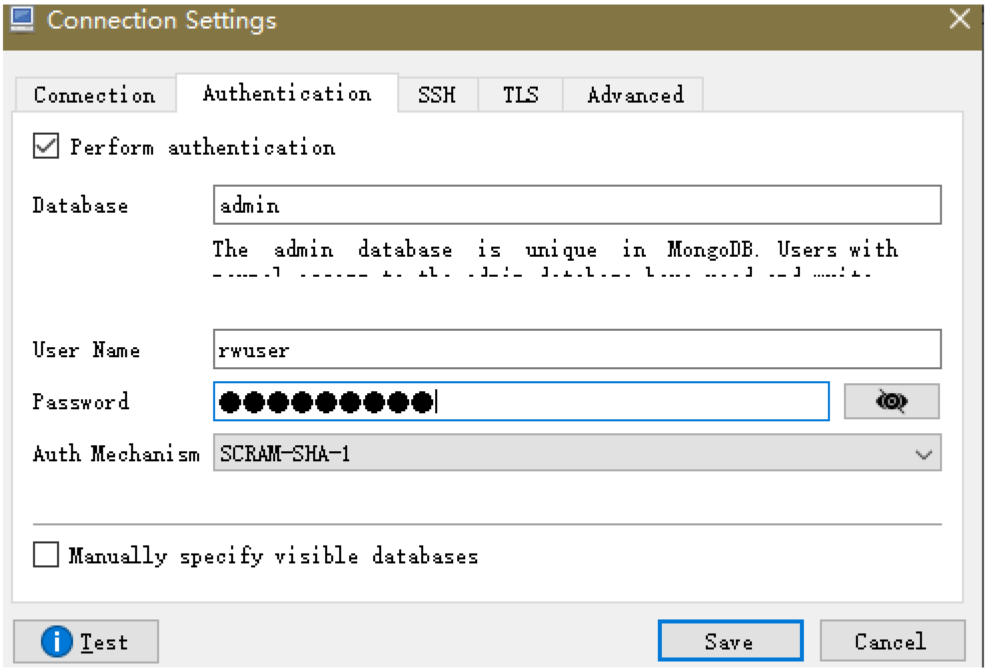

Connection Settings

Connection | Authentication | SSH | TLS | Advanced

☑ Perform authentication

Database: admin

The admin database is unique in MongoDB. Users with

User Name: rwuser

Password: ●●●●●●●●●●

Auth Mechanism: SCRAM-SHA-1

☐ Manually specify visible databases

3. On the **TLS** tab, select **Use TLS protocol** and select **Self-signed Certificate** for **Authentication Method**.

Figure 2-116 SSL

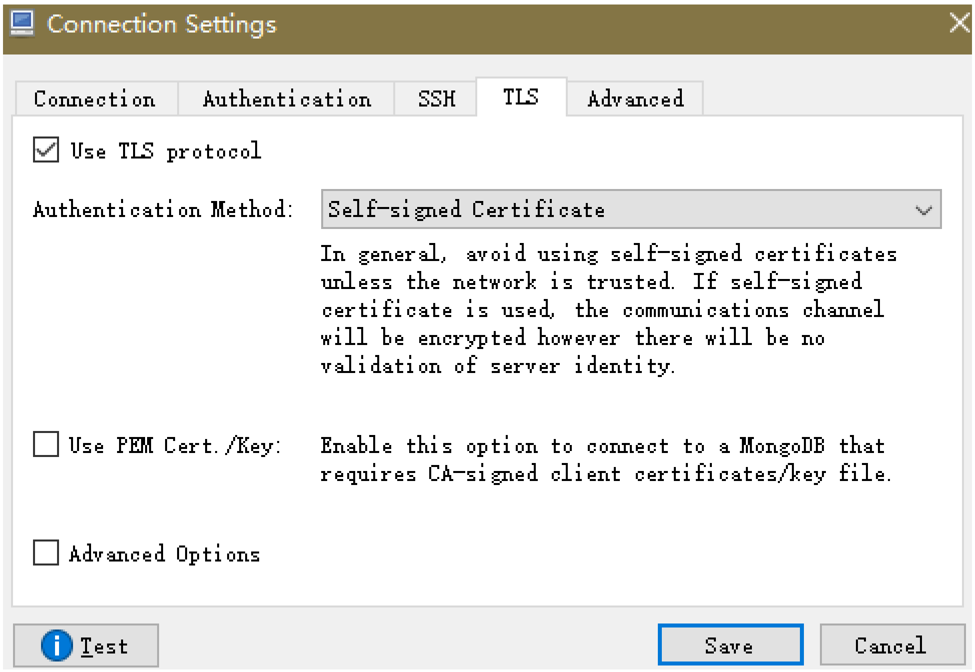

☑ Use TLS protocol

Authentication Method: Self-signed Certificate

In general, avoid using self-signed certificates unless the network is trusted. If self-signed certificate is used, the communications channel will be encrypted however there will be no validation of server identity.

□ Use PEM Cert./Key: Enable this option to connect to a MongoDB that requires CA-signed client certificates/key file.

□ Advanced Options

4. Click **Save**.

Step 3 On the **MongoDB Connections** page, click **Connect** to connect to the single node instance.

Figure 2-117 Single node connection information

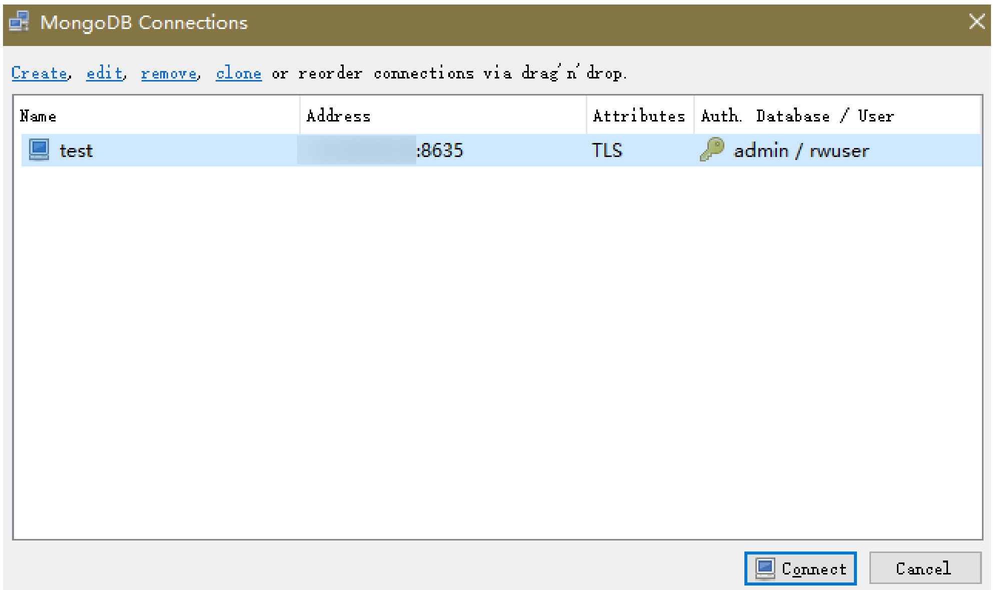

Create, edit, remove, clone or reorder connections via drag'n'drop.

Name	Address	Attributes	Auth. Database / User
test	[redacted]:8635	TLS	admin / rwuser

Step 4 If the single node instance is successfully connected, the page shown in **Figure 2-118** is displayed.

Figure 2-118 Single node connected

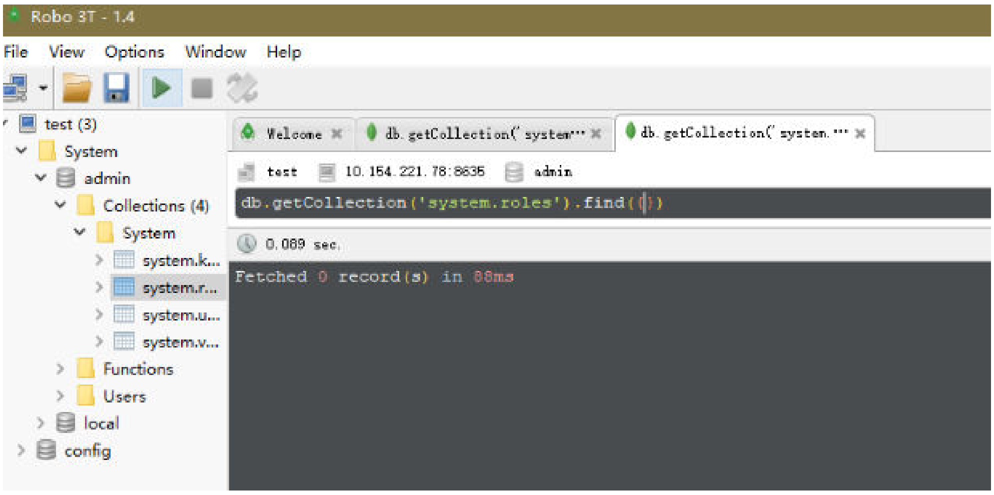

```
db.getCollection('system.roles').find({})
```

```
Fetched 0 record(s) in 88ms
```

----End

2.3.5 Connecting to a Single Node Instance Using Program Code

2.3.5.1 Java

If you are connecting to an instance using Java, an SSL certificate is optional, but downloading an SSL certificate and encrypting the connection will improve the security of your instance. SSL is disabled by default for newly created DB instances. You can enable SSL by referring to **Enabling or Disabling SSL**. SSL encrypts connections to databases but it increases the connection response time and CPU usage. Therefore, you are advised not to enable SSL.

Prerequisites

Familiarize yourself with:

- Computer basics
- Java code

Obtaining and Using Java

- Download the Jar driver from: **https://repo1.maven.org/maven2/org/mongodb/mongo-java-driver/3.0.4/**
- To view the usage guide, visit **https://mongodb.github.io/mongo-java-driver/4.2/driver/getting-started/installation/**.

Using an SSL Certificate

NOTE

- Download the SSL certificate and verify the certificate before connecting to databases.
- On the **Instances** page, click the target DB instance name. In the **DB Information** area on the **Basic Information** page, click ⤓ in the **SSL** field to download the root certificate or certificate bundle.
- For details about how to set up an SSL connection, see the MongoDB Java Driver official document at **https://www.mongodb.com/docs/drivers/java/sync/current/fundamentals/connection/tls/#std-label-tls-ssl**.
- Java Runtime Environment (JRE) earlier than Java 8 enables TLS 1.2 only in updated versions. If TLS 1.2 is not enabled for your JRE, upgrade it to a later version to use TLS 1.2 for connection.

Connect to a single node instance using Java. The format of the Java link is as follows:

```
mongodb://<username>:<password>@<instance_ip>:<instance_port>/<database_name>?authSource=admin&ssl=true
```

Table 2-32 Parameter description

Parameter	Description
<username>	Current username.
<password>	Password for the current username
<instance_ip>	If you attempt to access the instance from an ECS, set *instance_ip* to the private IP address displayed on the **Basic Information** page of the instance to which you intend to connect.
	If you intend to access the instance through an EIP, set *instance_ip* to the EIP that has been bound to the instance.
<instance_port>	Database port displayed on the **Basic Information** page. Default value: **8635**
<database_name>	Name of the database to be connected.
authSource	Authentication user database. The value is **admin**.
ssl	Connection mode. **true** indicates that the SSL connection mode is used.

Use the keytool to configure the CA certificate. For details about the parameters, see **Table 2-33**.

```
keytool -importcert -trustcacerts -file <path to certificate authority file> -keystore <path to trust store> -storepass <password>
```

Table 2-33 Parameter description

Parameter	Description
<path to certificate authority file>	Path for storing the SSL certificate.
<path to trust store>	Path for storing the truststore. Set this parameter as required, for example, **./trust/certs.keystore**.
<password>	Custom password.

Set the JVM system properties in the program to point to the correct truststore and keystore:

- System.setProperty("javax.net.ssl.trustStore","<path to trust store>");
- System.setProperty("javax.net.ssl.trustStorePassword","<password>");

For details about the Java code, see the following example:

```
public class Connector {
    public static void main(String[] args) {
        try {
            System.setProperty("javax.net.ssl.trustStore", "./trust/certs.keystore");
            System.setProperty("javax.net.ssl.trustStorePassword", "123456");
            ConnectionString connString = new ConnectionString("mongodb://
<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin&ssl=true");
            MongoClientSettings settings = MongoClientSettings.builder()
                    .applyConnectionString(connString)
                    .applyToSslSettings(builder -> builder.enabled(true))
                    .applyToSslSettings(builder -> builder.invalidHostNameAllowed(true))
                    .build();
            MongoClient mongoClient = MongoClients.create(settings);
            MongoDatabase database = mongoClient.getDatabase("admin");
            //Ping the database. If the operation fails, an exception occurs.
            BsonDocument command = new BsonDocument("ping", new BsonInt64(1));
            Document commandResult = database.runCommand(command);
            System.out.println("Connect to database successfully");
        } catch (Exception e) {
            e.printStackTrace();
            System.out.println("Test failed");
        }
    }
}
```

Connection Without the SSL Certificate

NOTE

You do not need to download the SSL certificate because certificate verification on the server is not required.

Connect a single node using Java. The Java link format is as follows:

```
mongodb://<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin
```

Table 2-34 Parameter description

Parameter	Description
<username>	Current username.
<password>	Password for the current username
<instance_ip>	If you attempt to access the instance from an ECS, set *instance_ip* to the private IP address displayed on the **Basic Information** page of the instance to which you intend to connect.
	If you intend to access the instance through an EIP, set *instance_ip* to the EIP that has been bound to the instance.
<instance_port>	Database port displayed on the **Basic Information** page. Default value: **8635**
<database_name>	Name of the database to be connected.
authSource	Authentication user database. The value is **admin**.

Example script in Java:

```
public class Connector {
    public static void main(String[] args) {
        try {
            ConnectionString connString = new ConnectionString("mongodb://
<username>:<password>@<instance_ip>:<instance_port>/<database_name>?
authSource=admin");
            MongoClientSettings settings = MongoClientSettings.builder()
                    .applyConnectionString(connString)
                    .retryWrites(true)
                    .build();
            MongoClient mongoClient = MongoClients.create(settings);
            MongoDatabase database = mongoClient.getDatabase("admin");
            //Ping the database. If the operation fails, an exception occurs.
            BsonDocument command = new BsonDocument("ping", new BsonInt64(1));
            Document commandResult = database.runCommand(command);
            System.out.println("Connect to database successfully");
        } catch (Exception e) {
            e.printStackTrace();
            System.out.println("Test failed");
        }
    }
}
```

2.3.5.2 Python

This section describes how to connect to a single node instance using Python.

Prerequisites

1. To connect an ECS to an instance, the ECS must be able to communicate with the DDS instance. You can run the following command to connect to the IP address and port of the instance server to test the network connectivity.

curl *ip:port*

If the message **It looks like you are trying to access MongoDB over HTTP on the native driver port** is displayed, the network connectivity is normal.

2. Install Python and third-party installation package **pymongo** on the ECS. Pymongo 2.8 is recommended.
3. If SSL is enabled, you need to download the root certificate and upload it to the ECS.

Connection Code

- Enabling SSL

```
import ssl
from pymongo import MongoClient
conn_urls="mongodb://rwuser:rwuserpassword@ip:port/{mydb}?authSource=admin"
connection = MongoClient(conn_urls,connectTimeoutMS=5000,ssl=True,
ssl_cert_reqs=ssl.CERT_REQUIRED,ssl_match_hostname=False,ssl_ca_certs=${path to
certificate authority file})
dbs = connection.database_names()
print "connect database success! database names is %s" % dbs
```

- Disabling SSL

```
import ssl
from pymongo import MongoClient
conn_urls="mongodb://rwuser:rwuserpassword@ip:port/{mydb}?authSource=admin"
connection = MongoClient(conn_urls,connectTimeoutMS=5000)
dbs = connection.database_names()
print "connect database success! database names is %s" % dbs
```

NOTE

- The authentication database in the URL must be **admin**. That means setting **authSource** to **admin**.
- In SSL mode, you need to manually generate the trustStore file.
- The authentication database must be **admin**, and then switch to the service database.

2.3.5.3 PHP

This section describes how to connect to a single node instance using PHP.

Prerequisites

1. To connect an ECS to a DDS instance, run the following command to connect to the IP address and port of the instance server to test the network connectivity.

   curl ip:port

   If the message **It looks like you are trying to access MongoDB over HTTP on the native driver port** is displayed, the ECS and DDS instance can communicate with each other.
2. If SSL is enabled, you need to download the root certificate and upload it to the ECS.

Obtaining and Using PHP

For the information about PHP, visit **https://www.php.net/mongodb-driver-manager.construct**

Connection Code

- Enabling SSL
  - Run **MongoDB\Client::__construct()** to create a client instance.

```
function __construct(
    ?string $uri = null,
    array $uriOptions = [],
    array $driverOptions = []
)
```

  - Use $uriOptions to set **SSL** to **true** to enable the SSL connection. Use $driverOptions to set **ca_file** to the CA certificate path and **allow_invalid_hostname** to **true**.

```
<?php

require 'vendor/autoload.php'; // include Composer goodies

$replicaset_url = 'mongodb://rwuser:*****@192.168.***.***:8635/test?
authSource=admin';
$test_db = 'test_db';
$test_coll = 'test_coll';

//Create mongoclient.
$client = new MongoDB\Client(
....$replicaset_url,
    [
        'ssl' => true,
    ],
    [
        "ca_file" => "/path/to/ca.pem",
        "allow_invalid_hostname" => true
    ]
);

$collection = $client->$test_db->$test_coll;

//Insert a record.
$result = $collection->insertOne([
    'username' => 'admin',
    'email' => 'admin@example.com',
]);
echo "Object ID: '{$result->getInsertedId()}'", "\n";

//Query a record.
$result = $collection->find(['username' => 'admin']);
foreach ($result as $entry) {
    echo $entry->_id, ': ', $entry->email, "\n";
}

?>
```

- Disabling SSL

```
<?php

require 'vendor/autoload.php'; // include Composer goodies

$replicaset_url = 'mongodb://rwuser:*****@192.168.***.***:8635/test?authSource=admin';
$test_db = 'test_db';
$test_coll = 'test_coll';

//Create mongoclient.
$client = new MongoDB\Client($replicaset_url);
```

```
$collection = $client->$test_db->$test_coll;

//Insert a record.
$result = $collection->insertOne([
    'username' => 'admin',
    'email' => 'admin@example.com',
]);
echo "Object ID: '{$result->getInsertedId()}'", "\n";

//Query a record.
$result = $collection->find(['username' => 'admin']);
foreach ($result as $entry) {
    echo $entry->_id, ': ', $entry->email, "\n";
}

?>
```

NOTE

- The authentication database in the URL must be **admin**. Set **authSource** to **admin**.
- The authentication database of the **rwuser** user must be **admin**.

3 Data Migration

3.1 Migration Scheme Overview

DDS provides multiple migration schemes to migrate MongoDB databases in different service scenarios.

Table 3-1 Migration schemes

Scenario	Migration Types	References
Migrating data using the export and import tools	Full	• **Migrating Data Using mongoexport and mongoimport** • **Migrating Data Using mongodump and mongorestore**
Migrating data from other cloud MongoDB to DDS	Full +incremental	**Migrating from Other Cloud MongoDB to DDS**
Migrating data from on-premises MongoDB to DDS	Full +incremental	**Migrating from On-Premises MongoDB to DDS**
Migrating data from ECS-hosted MongoDB to DDS	Full +incremental	**Migrating from ECS MongoDB Databases to DDS**
Migrating data from DDS to MongoDB	Full +incremental	**Migrating from DDS to MongoDB**

3.2 Migrating Data Using DRS

Data Replication Service (DRS) helps migrate your databases to DDS DB instances. During the migration, the source remains operational even if a transfer is interrupted, thereby minimizing application downtime.

Prerequisites

To improve the stability and security of your migration ensure that your instances meet the migration requirements described in **Migration Preparations**.

Migration Types

- **Full migration**

  This migration type is suitable for scenarios where some service interruptions are acceptable. All objects and data in non-system databases are migrated to the destination database in a single batch. The objects include tables, views, and stored procedures. If you perform a full migration, stop operations on the source database, or data generated in the source database during the migration will result in inconsistencies with the destination database.

- **Full+Incremental migration**

  This migration type allows you to migrate data without interrupting services. After a full migration initializes the destination database, an incremental migration initiates and parses logs to ensure data consistency between the source and destination databases. If you select the **Full+Incremental** migration type, data generated during the full migration will be synchronized to the destination database with zero downtime, ensuring that both the source and destination databases remain accessible throughout the process.

Supported Source and Destination Databases

Table 3-2 Supported databases

Source DB	Destination DB
● On-premises Mongo (versions 3.2, 3.4, and 4.0) ● Self-built MongoDB on ECSs (versions 3.2, 3.4, and 4.0) ● MongoDB 3.2, 3.4, and 4.0 on other clouds (Tencent Cloud MongoDB 3.2 is not supported.) ● DDS DB instances (versions 3.4 and 4.0)	● DDS DB instances (versions 3.4, 4.0, and 4.2) **NOTE** The destination database version must be the same as or later than the source database version.

Supported Migration Objects

Different types of migration tasks support different migration objects. For details, see **Table 3-3**. DRS will automatically check the objects you selected before the migration.

Table 3-3 Migration objects

Type	Precautions
Migration objects	● Object level: table level, database level, or instance level (full migration). ● Supported migration objects: – Associated objects must be migrated at the same time to avoid migration failure caused by missing associated objects. Common associations: collections referenced by views, and views referenced by views – Replica set: Only collections (including validator and capped collections), indexes, and views can be migrated. – Cluster: Only collections (including validator and capped collections), shard keys, indexes, and views can be migrated. – Single node: Only collections (including validator and capped collections), indexes, and views can be migrated. – Only user data and source database account information can be migrated. The system databases (for example, local, admin, and config) and system collection cannot be migrated. If service data is stored in a system database, run the **renameCollection** command to move the service data to the user database. – The statement for creating a view cannot contain a regular expression. – Collections that contain the **_id** field without indexes are not supported. – The first parameter of **BinData()** cannot be **2**. – If ranged sharding is used, **maxKey** cannot be used as the primary key. **NOTE** The objects that can be migrated have the following constraints: ● The source database name cannot contain /\."$ or spaces. The collection name and view name cannot start with **system.** or contain the dollar sign ($).

Database Account Permission Requirements

To start a migration task, the source and destination database users must have permissions listed in the following table. Different types of migration tasks require different permissions. For details, see **Table 3-4**. DRS automatically checks the database account permissions in the pre-check phase and provides handling suggestions.

NOTE

- You are advised to create an independent database account for DRS task connection to prevent task failures caused by database account password modification.
- After changing the account passwords for the source and destination databases, **modify the connection information** in the DRS task as soon as possible to prevent automatic retry after a task failure. Automatic retry will lock the database accounts.

Table 3-4 Database account permission

Type	Full migration	Full+Incremental Migration
Source database user	• Replica set: The source database user must have the readAnyDatabase permission for the admin database. • Single node: The source database user must have the readAnyDatabase permission for the admin database. • Cluster: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database. • To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the **system.users** and **system.roles** system tables of the admin database.	• Replica set: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database. • Single node: The source database user must have the readAnyDatabase permission for the admin database and the read permission for the local database. • Cluster: The source dds mongos node user must have the readAnyDatabase permission for the admin database and the read permission for the config database. The source shard node user must have the readAnyDatabase permission for the admin database and the read permission for the local database. • To migrate accounts and roles of the source database, the source and destination database users must have the read permission for the **system.users** and **system.roles** system tables of the admin database.

Type	Full migration	Full+Incremental Migration
Destination database user	The destination database user must have the dbAdminAnyDatabase permission for the admin database and the readWrite permission for the destination database. If the destination database is a cluster instance, the database user must have the clusterManager permission for the admin database.	

NOTE

For example, the source database user must have the readAnyDatabase permission for the admin database and the read permission for the config database.

```
db.grantRolesToUser("Username",[{role:"readAnyDatabase",db:"admin"},
{role:"read",db:"config"}])
```

Migration Operations

For details, see **MongoDB Database Migration** in *Data Replication Service Best Practices*.

3.3 Migrating Data Using mongoexport and mongoimport

mongoexport and mongoimport are backup and restoration tools provided by the MongoDB client. You can install a MongoDB client on the local device or ECS and use the mongoexport and mongoimport tools to migrate your on-premises MongoDB databases or other cloud MongoDB databases to DDS instances.

Before migrating data from a MongoDB database to DDS, transfer data to a .json file using the mongoexport tool. This section describes how to import the data from the JSON files to DDS using the mongoimport tool on the ECS or from some other devices that can access DDS.

Precautions

- The mongoexport and mongoimport tools support only full migration. To ensure data consistency, stop services on the source database and stop writing data to the source database before the migration.
- You are advised to perform the migration during off-peak hours to avoid impacting services.
- The admin and local system databases cannot be migrated.
- Make sure that no service set has been created in the system databases admin and local in the source database. If there is already a service set, migrate them out of the system databases admin and local before migration.
- Before importing data, ensure that the necessary indexes are there on the source database. Delete any unnecessary indexes and create any necessary indexes before migration.

- If you choose to migrate a sharded cluster, you must create a set of shards in the destination database and configure sharding. In addition, indexes must be created before migration.

Prerequisites

1. An ECS or a device that can access DDS is ready for use.
   - To connect to a DDS DB instance through a private network from an ECS, create and log in to the ECS. For details, see **Purchasing an ECS** and **Logging In to an ECS**.
   - To bind an EIP to a DB instance:
     i. Bind an EIP to a node in the instance. For details about how to bind an EIP to a node, see "Binding an EIP" in *Getting Started with Document Database Service*.
     ii. Ensure that your local device can access the EIP that has been bound to the DB instance.
2. A migration tool has been installed on the prepared ECS.

   For details on how to install the migration tool, see **How Can I Install a MongoDB Client?**

   NOTE

   The MongoDB client provides the mongoexport and mongoimport tools.

Exporting Data

Step 1 Log in to the ECS or the device that can access DDS.

Step 2 Use the mongoexport tool to transfer data from the source database to a .json file.

The SSL connection is used as an example. If you select a common connection, delete **--ssl --sslAllowInvalidCertificates** from the following command.

./**mongoexport --host** *<DB_ADDRESS>* **--port** *<DB_PORT>* **--ssl --sslAllowInvalidCertificates --type json --authenticationDatabase** *<AUTH_DB>* **-u** *<DB_USER>* **--db** *<DB_NAME>* **--collection** *<DB_COLLECTION>* **--out** *<DB_PATH>*

- **DB_ADDRESS** is the database address.
- **DB_PORT** is the database port.
- **AUTH_DB** is the database for storing DB_USER information. Generally, this value is **admin**.
- **DB_USER** is the database user.
- **DB_NAME** is the name of the database from which data will be exported.
- **DB_COLLECTION** is the collection of the database from which data will be exported.
- **DB_PATH** is the path where the .json file is located.

Enter the database administrator password when prompted:

```
Enter password:
```

The following is an example. After the command is executed, the **exportfile.json** file will be generated:

./mongoexport --host 192.168.1.21 --port 8635 --ssl --sslAllowInvalidCertificates --type json --authenticationDatabase admin -u rwuser --db test02 --collection Test --out /tmp/mongodb/export/exportfile.json

Step 3 View the results

If information similar to the following is displayed, the data has been successfully exported. **x** is the number of exported data records.

```
exported x records
```

Step 4 Compress the exported .json file.

gzip exportfile.json

Compressing the file helps reduce the time needed to transmit the data. The compressed file is **exportfile.json.gz**.

----End

Importing Data

Step 1 Log in to the ECS or whichever device you will be using to access DDS.

Step 2 Upload the data to be imported to the ECS or the device.

Select an uploading method based on the OS you are using.

- In Linux, for example, you can use secure copy protocol (SCP):

  scp *<IDENTITY_FILE>*
  <REMOTE_USER>@*<REMOTE_ADDRESS>*:*<REMOTE_DIR>*

  - **IDENTITY_FILE** is the directory where the **exportfile.json.gz** file is located. The file access permission is 600.
  - **REMOTE_USER** is the ECS OS user.
  - **REMOTE_ADDRESS** is the ECS address.
  - **REMOTE_DIR** is the directory of the ECS to which the **exportfile.json.gz** file is uploaded.
- In Windows, upload **exportfile.json.gz** to the ECS using file transfer tools.

Step 3 Decompress the package.

gzip -d *exportfile.json.gz*

Step 4 Import the JSON file to the DDS database.

The SSL connection is used as an example. If you select a common connection, delete **--ssl --sslAllowInvalidCertificates** from the following command.

./**mongoimport --host** *<DB_ADDRESS>* **--port** *<DB_PORT>* **--ssl --sslAllowInvalidCertificates --type json --authenticationDatabase** *<AUTH_DB>* **-u** *<DB_USER>* **--db** *<DB_NAME>* **--collection** *<DB_COLLECTION>* **--file** *<DB_PATH>*

- **DB_ADDRESS** indicates the DB instance IP address.
- **DB_PORT** indicates the database port.
- **AUTH_DB** indicates the database that authenticates DB_USER. Generally, this value is **admin**.
- **DB_USER** indicates the account name of the database administrator.
- **DB_NAME** indicates the name of the database to which data will be imported.
- **DB_COLLECTION** indicates the collection of the database to which data will be imported.
- **DB_PATH** indicates the path where the .json file is located.

Enter the database administrator password when prompted:

```
Enter password:
```

The following is an example:

./mongoimport --host 192.168.1.21 --port 8635 --ssl --sslAllowInvalidCertificates --type json --authenticationDatabase admin -u rwuser --db test02 --collection Test --file /tmp/mongodb/export/exportfile.json

Step 5 View the results.

If information similar to the following is displayed, the data has been successfully imported. **x** is the number of imported data records.

```
imported x records
```

----End

3.4 Migrating Data Using mongodump and mongorestore

mongodump and mongorestore are backup and restoration tools provided by the MongoDB client. You can install a MongoDB client on the local device or ECS and use the mongodump and mongorestore tools to migrate your on-premises MongoDB databases or other cloud MongoDB databases to DDS instances.

Precautions

- The mongodump and mongorestore tools support only full migration. To ensure data consistency, stop services on the source database and stop writing data to the source database before the migration.
- You are advised to perform the migration during off-peak hours to avoid impacting services.
- The admin and local system databases cannot be migrated.
- The file exported by mongodump is a BSON binary file. The mongorestore uses this binary backup file to restore data to a DB instance.
- Make sure that no service set has been created in the system databases admin and local in the source database. If there is already a service set, migrate them out of the system databases admin and local before migration.

- Before importing data, ensure that the necessary indexes are there on the source database. Delete any unnecessary indexes and create any necessary indexes before migration.
- If you choose to migrate a sharded cluster, you must create a set of shards in the destination database and configure sharding. In addition, indexes must be created before migration.
- If the backup using the mongodump tool fails (for example, an error is reported when the backup progress reaches 97%), you are advised to increase the storage space of the VM that fails to be backed up and reserve some redundant space before performing the backup again.
- User **rwuser** can only operate service database tables. You are advised to specify databases and tables to import and export only service data. Otherwise, the insufficient permission problem may occur during full import and export.
- For details about how to restore backup data to an on-premises database, see **Restoring Data to an On-Premises Database**.

Prerequisites

1. Prepare an ECS or a device that can access DDS.
   - To connect to a DDS instance over a private network from an ECS, create and log in to the ECS. For details, see **Purchasing an ECS** and **Logging In to an ECS**.
   - To bind an EIP to a DB instance:
     i. Bind an EIP to a node in the DB instance. For details about how to bind an EIP to a node, see "Binding an EIP" in the *Getting Started with Document Database Service* .
     ii. Ensure that your local device can access the EIP that has been bound to the DB instance.
2. A migration tool has been installed on the prepared ECS.

   For details on how to install the migration tool, see **How Can I Install a MongoDB Client?**

   NOTE

   - The mongodump and mongorestore tools are part of the MongoDB client installation package.
   - The MongoDB client version must match the instance version. Otherwise, compatibility issues may occur.

Exporting Data

Step 1 Log in to the ECS or the device that can access DDS.

Step 2 Back up the source database data using the mongodump tool.

An SSL connection is used in this example. If you select an unencrypted connection, delete **--ssl --sslCAFile** <*FILE_PATH*> **--sslAllowInvalidCertificates** from the following command.

./mongodump --host *<DB_HOST>* **--port** *<DB_PORT>* **--authenticationDatabase** *<AUTH_DB>* **-u <DB_USER> --ssl --sslCAFile** <*FILE_PATH*> **--**

sslAllowInvalidCertificates --db <*DB_NAME*> **--collection** <*DB_COLLECTION*> **--gzip --archive=**<*Name of the backup file that contains the file path*>

Table 3-5 Parameter description

Parameter	Description
<DB_HOST>	Database address
<DB_PORT>	Database port
<DB_USER>	Database username
<AUTH_DB>	Database that stores *<DB_USER>* information. Generally, the value is **admin**.
<*FILE_PATH*>	Path for storing the root certificate
<*DB_NAME*>	The name of the database to be migrated.
<*DB_COLLECTION*>	Collection in the database to be migrated

Enter the database administrator password when prompted:

```
Enter password:
```

After the command is executed, the file specified by **archive** is the final backup file. The following command uses **backup.tar.gz** as an example.

./mongodump --host 192.168.*xx.xx* **--port 8635 --authenticationDatabase admin -u rwuser --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidCertificates --db test --collection usertable --gzip --archive=backup.tar.gz**

```
2019-03-04T18:42:10.687+0800    writing admin.system.users to
2019-03-04T18:42:10.688+0800    done dumping admin.system.users (1 document)
2019-03-04T18:42:10.688+0800    writing admin.system.roles to
2019-03-04T18:42:10.690+0800    done dumping admin.system.roles (0 documents)
2019-03-04T18:42:10.690+0800    writing admin.system.version to
2019-03-04T18:42:10.691+0800    done dumping admin.system.version (2 documents)
2019-03-04T18:42:10.691+0800    writing test.test_collection to
2019-03-04T18:42:10.691+0800    writing admin.system.profile to
2019-03-04T18:42:10.692+0800    done dumping admin.system.profile (4 documents)
2019-03-04T18:42:10.695+0800    done dumping test.test_collection (198 documents)
```

----End

Importing Data

Step 1 Log in to the ECS or whichever device you will be using to access DDS.

Step 2 Upload the data to be imported to the ECS or the device.

Select an uploading method based on the OS you are using.

- In Linux, for example, you can use secure copy protocol (SCP):

  scp -r *<IDENTITY_DIR> <REMOTE_USER>*@*<REMOTE_ADDRESS>*:*<REMOTE_DIR>*

Table 3-6 Parameter description

Parameter	Description
<IDENTITY_DIR>	Directory where the backup folder is located.
<REMOTE_USER>	User of ECS OS in **Step 1**
<REMOTE_ADDRESS>	IP address of the ECS in **Step 1**
<REMOTE_DIR>	Directory of the ECS to be imported

- In Windows, upload the backup directory to the ECS using a file transfer tool.

Step 3 Import the backup data to DDS.

An SSL connection is used in this example. If you use an unencrypted connection, delete **--ssl --sslCAFile** <*FILE_PATH*> **--sslAllowInvalidCertificates** from the following command.

./mongorestore --host *<DB_HOST>* **--port** *<DB_PORT>* **--authenticationDatabase** *<AUTH_DB>* **-u***<DB_USER>* **--ssl --sslCAFile** <*FILE_PATH*> **--sslAllowInvalidCertificates --db** <*DB_NAME*> **--collection** <*DB_COLLECTION*> **--gzip --archive=**<*Name of the backup file that contains the file path*>

Table 3-7 Parameter description

Parameter	Description
<DB_HOST>	DDS database address
<DB_PORT>	Database port
<AUTH_DB>	The database that authenticates *DB_USER*. Generally, the value is **admin**.
<DB_USER>	Account name of the database administrator. The default value is **rwuser**.
<*FILE_PATH*>	Path for storing the root certificate
<*DB_NAME*>	The name of the database to be migrated.
<*DB_COLLECTION*>	Collection in the database to be migrated

Enter the database administrator password when prompted:

```
Enter password:
```

The following is an example:

./mongorestore --host 192.168.*xx.xx* **--port 8635 --authenticationDatabase admin -u rwuser --ssl --sslCAFile /tmp/ca.crt --sslAllowInvalidCertificates --db test --collection usertable --gzip --archive=backup.tar.gz**

```
2019-03-05T14:19:43.240+0800    preparing collections to restore from
2019-03-05T14:19:43.243+0800    reading metadata for test.test_collection from dump/test/
test_collection.metadata.json
2019-03-05T14:19:43.263+0800    restoring test.test_collection from dump/test/test_collection.bson
2019-03-05T14:19:43.271+0800    restoring indexes for collection test.test_collection from metadata
2019-03-05T14:19:43.273+0800    finished restoring test.test_collection (198 documents)
2019-03-05T14:19:43.273+0800    restoring users from dump/admin/system.users.bson
2019-03-05T14:19:43.305+0800    roles file 'dump/admin/system.roles.bson' is empty; skipping roles
restoration
2019-03-05T14:19:43.305+0800    restoring roles from dump/admin/system.roles.bson
2019-03-05T14:19:43.333+0800    done
```

----End

Related Issues

When you back up the entire instance using mongodump and mongorestore, the permission verification fails.

- Cause

  The **rwuser** user has limited permissions on the **admin** and **config** databases of the instance. As a result, the permission verification fails.

- Solution

  Grant permissions on certain databases and tables to the user.

4 Performance Tuning

4.1 Parameters

Database parameters are key configuration items in a database system. Improper parameter settings may adversely affect database performance. This document describes some important parameters. For details on parameter descriptions, visit **MongoDB official website**.

For details about how to change parameter values on the console, see **Modifying DDS DB Instance Parameters**.

- **enableMajorityReadConcern**

  This parameter indicates whether data read has been acknowledged by a majority of nodes.

  The default value is **false**, indicating that data read is returned after being acknowledged by a single node.

  If this parameter is set to **true**, data read is returned after being acknowledged by a majority of nodes. This operation will increase the size of the LAS file, resulting in high CPU usage and disk usage.

  In DDS, read concern cannot be set to majority. If majority read concern is required, you can set write concern to majority, indicating that data is written to a majority of nodes. In this way, data on most nodes is consistent. Then, by reading data from a single node, it can be ensured that the data has been written to a majority of nodes, and there are no dirty reads.

  NOTE

  Write concern and read concern respectively specify the write and read policies for MongoDB.

  If read concern is set to majority, data read by users has been written to a majority of nodes and will not be rolled back to avoid dirty reads.

- **failIndexKeyTooLong**

  The default value is **true**.

  This parameter cannot be modified to avoid an excessively long index key.

- **net.maxIncomingConnections**

This parameter indicates the maximum number of concurrent connections that dds mongos or mongod can accept. The default value depends on the **instance specifications**. This parameter is displayed as **default** before being set, indicating that the parameter value varies with the memory specifications.

- **security.javascriptEnabled**

  The default value is **false**.

  This parameter indicates whether JavaScript scripts can be executed on mongod. For security purposes, the default value is **false**, indicating that JavaScript scripts cannot be executed on mongod, and the mapreduce and group commands cannot be used.

- **disableJavaScriptJIT**

  The default value is **true**.

  This parameter indicates whether to disable JavaScript JIT compilation. JavaScript JIT compilation enables just-in-time (JIT) compilation to improve the performance of running scripts.

  disableJavaScriptJIT: The default value is **true**, indicating that the JavaScriptJIT compiler is disabled. To enable JavaScript JIT compilation, set **disableJavaScriptJIT** to **false**.

- **operationProfiling.mode**

  The parameter value is **slowOp** by default.

  This parameter indicates the level of the database analyzer.

  This parameter supports the following values:

  - The default value is **slowOp**, indicating that the collector records statements whose response time exceeds the threshold.
  - The value **off** indicates that the analyzer is disabled and does not collect any data.
  - The value **all** indicates that the collector collects data of all operations.

- **operationProfiling.slowOpThresholdMs**

  The default value is **500** and the unit is ms.

  This parameter indicates the threshold for slow queries in the unit of ms. Queries that take longer than the threshold are deemed as slow queries.

  Unless otherwise specified, setting the value to 500 ms is recommended.

- **maxTransactionLockRequestTimeoutMillis**

  The value ranges from **5** to **100**, in milliseconds. The default value is **5**.

  This parameter specifies the time for a transaction to wait for locks. If the time is exceeded, the transaction is rolled back.

4.2 Read and Write Performance

Common check items:

1. If the error message Timeout is displayed in the database, check whether the number of connections to the instance reaches the upper limit.
   - Check method: View the **monitoring metric** to check whether the **maximum number of active connections** has been reached.

   - Solution: See **What Can I Do If the Number of Connections of an Instance Reaches Its Maximum?**

2. Check whether the instance is properly connected.
   - Check method: Check whether multiple dds mongos nodes in a cluster instance are connected and whether both the primary and secondary nodes in a replica set instance are connected.
   - Solution: If you connect to a cluster instance, connect to multiple dds mongos nodes at the same time to share the load and improve availability. If you connect to a replica set instance, connect to both the primary and secondary nodes. This improves read/write performance and prevents errors reported when data is written from the client after a primary/standby switchover.
3. Check whether the monitoring metrics of the instance are normal.
   - Check method: View **monitoring metrics** to check the CPU usage and memory usage.
   - Solution: If the CPU and memory metrics are abnormal, check whether the client service load is too centralized or instance data is too intensive. If the client service load is too centralized, optimize the client architecture. If data is too intensive, shard data.
4. Check whether there are too many slow query logs.

   Check method: For details, see **Viewing Slow Query Logs**.

   Solution: For details, see **Slow Operation Optimization**.

Other precautions:

- During the query, select only the fields that need to be returned. When modifying data, modify only the fields that need to be modified. Do not directly store all modifications of the entire object. In this way, the network and processing loads are reduced.
- In the same service scenario, reduce the number of interactions with the database and query data at a time if possible.
- In a single instance, the total number of databases cannot exceed 200, and the total number of collections cannot exceed 500.
- Before bringing a service online, perform a load test to measure the performance of the database in peak hours.
- Do not execute a large number of concurrent transactions at the same time or leave a transaction uncommitted for a long time.
- Before the service is brought online, execute the query plan to check the query performance for all query types.
- Check the performance baseline of the instance specifications and analyze whether the current service requirements reach the upper limit.

4.3 High CPU Usage

If your CPU usage reaches 80%, a CPU bottleneck exists. In this case, data read and write are slow, affecting your services.

The following describes how to analyze current slow queries. After the analysis and optimization, query performance will be improved and indexes will be used more efficiently.

Analyzing Current Queries

1. Connect to an instance using Mongo Shell.

   To enable public access, see:

   - **Connecting to a Cluster Instance over a Public Network**
   - **Connecting to a Replica Set Instance over a Public Network**
   - **Connecting to a Single Node Instance over a Public Network**

   To access an instance over a private network, see:

   - **Connecting to a Cluster Instance over a Private Network**
   - **Connecting to a Replica Set Instance over a Private Network**
   - **Connecting to a Single Node Instance over a Private Network**

2. Run the following command to view the operations being performed on the database:

   db.currentOp()

   Command output:

```
{
        "raw" : {
                "shard0001" : {
                        "inprog" : [
                                {
                                        "desc" : "StatisticsCollector",
                                        "threadId" : "140323686905600",
                                        "active" : true,
                                        "opid" : 9037713,
                                        "op" : "none",
                                        "ns" : "",
                                        "query" : {

                                        },
                                        "numYields" : 0,
                                        "locks" : {

                                        },
                                        "waitingForLock" : false,
                                        "lockStats" : {

                                        }
                                },
                                {
                                        "desc" : "conn2607",
                                        "threadId" : "140323415066368",
                                        "connectionId" : 2607,
                                        "client" : "172.16.36.87:37804",
                                        "appName" : "MongoDB Shell",
                                        "active" : true,
                                        "opid" : 9039588,
                                        "secs_running" : 0,
                                        "microsecs_running" : NumberLong(63),
                                        "op" : "command",
                                        "ns" : "admin.",
                                        "query" : {
                                                "currentOp" : 1
                                   },
                                        "numYields" : 0,
                                        "locks" : {
                    },
                    "waitingForLock" : false,
                    "lockStats" : {

                    }
                }
            ],
            "ok" : 1
        },
  ...
}
```

NOTE

- **client**: IP address of the client that sends the request
- **opid**: unique operation ID
- **secs_running**: elapsed time for execution, in seconds. If the returned value of this field is too large, check whether the request is reasonable.
- **microsecs_running**: elapsed time for execution, in seconds. If the returned value of this field is too large, check whether the request is reasonable.
- **op**: operation type. The operations can be query, insert, update, delete, or command.
- **ns**: target collection
- For details, see the **db.currentOp()** command in **official document**.

3. Based on the command output, check whether there are requests that take a long time to process.

   If the CPU usage is low while services are being processed but then becomes high during just certain operations, analyze the requests that take a long time to execute.

   If an abnormal query is found, find the **opid** corresponding to the operation and run **db.killOp(*opid*)** to kill it.

Analyzing Slow Queries

Slow query profiling is enabled for DDS by default. The system automatically records any queries whose execution takes longer than 500 ms to the **system.profile** collection in the corresponding database. You can:

1. Connect to an instance using Mongo Shell.

   To access an instance from the Internet

   For details, see

   – **Connecting to a Cluster Instance over a Public Network**

   – **Connecting to a Replica Set Instance over a Public Network**

   – **Connecting to a Single Node Instance over a Public Network**

   To access an instance that is not publicly accessible

   For details, see

   – **Connecting to a Cluster Instance over a Private Network**

   – **Connecting to a Replica Set Instance over a Private Network**

   – **Connecting to a Single Node Instance over a Private Network**

2. Select a specific database (using the **test** database as an example):

   use test

3. Check whether slow SQL queries have been collected in **system.profile**.

   show collections;

   – If the command output includes **system.profile**, slow SQL queries have been generated. Go to the next step.
     ```
     mongos> show collections
     system.profile
     test
     ```
   – If the command output does not contain **system.profile**, no slow SQL queries have been generated, and slow query analysis is not required.
     ```
     mongos> show collections
     test
     ```
4. Check the slow query logs in the database.

   db.system.profile.find().pretty()

5. Analyze slow query logs to find the cause of the high CPU usage.

   The following is an example of a slow query log. The log shows a request that scanned the entire table, including 1,561,632 documents and without using a search index.

```
{
        "op" : "query",
        "ns" : "taiyiDatabase.taiyiTables$10002e",
        "query" : {
                "find" : "taiyiTables",
                "filter" : {
                        "filed19" : NumberLong("852605039766")
                },
                "shardVersion" : [
                        Timestamp(1, 1048673),
                        ObjectId("5da43185267ad9c374a72fd5")
                ],
                "chunkId" : "10002e"
        },
        "keysExamined" : 0,
        "docsExamined" : 1561632,
        "cursorExhausted" : true,
        "numYield" : 12335,
        "locks" : {
                "Global" : {
                        "acquireCount" : {
                                "r" : NumberLong(24672)
                        }
                },
                "Database" : {
                        "acquireCount" : {
                                "r" : NumberLong(12336)
                        }
                },
                "Collection" : {
                        "acquireCount" : {
                                "r" : NumberLong(12336)
                        }
                }
        },
        "nreturned" : 0,
        "responseLength" : 157,
        "protocol" : "op_command",
        "millis" : 44480,
        "planSummary" : "COLLSCAN",
        "execStats" : {
              "stage" :
"SHARDING_FILTER",
                [3/1955]
                "nReturned" : 0,
                "executionTimeMillisEstimate" : 43701,
```

```
        "works" : 1561634,
        "advanced" : 0,
        "needTime" : 1561633,
        "needYield" : 0,
        "saveState" : 12335,
        "restoreState" : 12335,
        "isEOF" : 1,
        "invalidates" : 0,
        "chunkSkips" : 0,
        "inputStage" : {
            "stage" : "COLLSCAN",
            "filter" : {
                "filed19" : {
                    "$eq" : NumberLong("852605039766")
                }
            },
            "nReturned" : 0,
            "executionTimeMillisEstimate" : 43590,
            "works" : 1561634,
            "advanced" : 0,
            "needTime" : 1561633,
            "needYield" : 0,
            "saveState" : 12335,
            "restoreState" : 12335,
            "isEOF" : 1,
            "invalidates" : 0,
            "direction" : "forward",
            "docsExamined" : 1561632
        }
    },
    "ts" : ISODate("2019-10-14T10:49:52.780Z"),
    "client" : "172.16.36.87",
    "appName" : "MongoDB Shell",
    "allUsers" : [
        {
            "user" : "__system",
            "db" : "local"
        }
    ],
   "user" : "__system@local"
}
```

The following stages can be causes for a slow query:

– **COLLSCAN** involves a full collection (full table) scan.

  When a request (such as query, update, and delete) requires a full table scan, a large amount of CPU resources are occupied. If you find **COLLSCAN** in the slow query log, a full table scan was performed and that occupy a lot of CPU resources.

  If such requests are frequent, create indexes for the fields to be queried.

– **docsExamined** involves a full collection (full table) scan.

  You can view the value of **docsExamined** to check the number of documents scanned. A larger value indicates a higher CPU usage.

– **IXSCAN** and **keysExamined** scan indexes.

  NOTE

  - An excessive number of indexes can affect the write and update performance.
  - If your application has more write operations, creating indexes may increase write latency.

  You can view the value of **keyExamined** to see how many indexes are scanned in a query. A larger value indicates a higher CPU usage.

If an index is not properly created or there are many matching results, the CPU usage does not decrease greatly and the execution speed is slow.

Example: For the data of a collection, the number of values of the **a** field is small (only **1** and **2**), but the **b** field has more values.

```
{ a: 1, b: 1 }
{ a: 1, b: 2 }
{ a: 1, b: 3 }
......
{ a: 1, b: 100000}
{ a: 2, b: 1 }
{ a: 2, b: 2 }
{ a: 2, b: 3 }
......
{ a: 1, y: 100000}
```

The following shows how to implement the {a: 1, b: 2} query.

```
db.createIndex({a: 1}): The query is not effective because the a field has too many same values.
db.createIndex({a: 1, b: 1}): The query is not effective because the a field has too many same values.
db.createIndex({b: 1}): The query is effective because the b field has a few same values.
db.createIndex({b: 1, a: 1}): The query is not effective because the a field has a few same values.
```

For the differences between {a: 1} and {b: 1, a: 1}, see the **official documents**.

– **SORT** and **hasSortStage** may involve sorting a large amount of data.

If the value of the **hasSortStage** parameter in the **system.profile** collection is **true**, the query request involves sorting. If the sorting cannot be implemented through indexes, the query results are sorted, and sorting is a CPU intensive operation. In this scenario, you need to create indexes for fields that are frequently sorted.

If the **system.profile** collection contains **SORT**, you can use indexing to improve sorting speed.

Other operations, such as index creation and aggregation (combinations of traversal, query, update, and sorting), also apply to the above mentioned scenarios because they are also CPU intensive operations. For more information about profiling, see **official documents**.

Analysis Capability

After the analysis and optimization of the requests that are being executed and slow requests, all requests use proper indexes, and the CPU usage becomes stable. If the CPU usage remains high after the analysis and troubleshooting, the current instance may have reached the performance bottleneck and cannot meet service requirements. In this case, you can perform the following operations to solve the problem:

1. View monitoring information to analyze instance resource usage. For details, see **Viewing Monitoring Metrics**.
2. Change the DDS instance class or add shard nodes.

4.4 High Storage Usage

If the storage usage of a DDS instance is too high or fully used, the instance becomes unavailable.

This section describes how to analyze and fix high storage usage.

Checking the Storage Usage

DDS provides the following two methods to check the storage usage of an instance:

1. Check the storage usage on the DDS console.

   You can log in to the DDS console and click the instance. On the **Basic Information** page, you can view the storage space of the instance in the **Storage Space** area.

   Figure 4-1 Checking the storage usage

   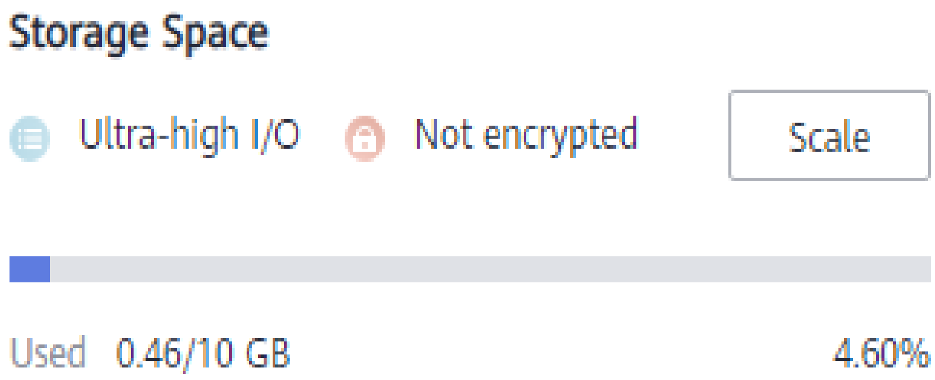

   Storage Space

   Ultra-high I/O Not encrypted

   Scale

   Used 0.46/10 GB 4.60%

2. View the monitoring metrics (storage usage and used storage).

   To view monitoring metrics, see **Viewing Monitoring Metrics**.

   Figure 4-2 Checking the storage usage

   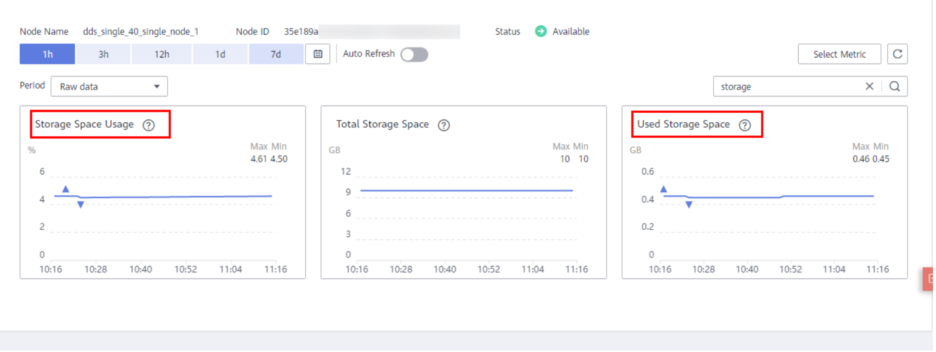

   Node Name dds_single_40_single_node_1 Node ID 35e189a[redacted] Status Available

   1h 3h 12h 1d 7d Auto Refresh Select Metric

   Period Raw data

   storage

Solution

1. For cluster instances, data may be unevenly distributed because the database collection is not properly sharded. As a result, the storage usage is high.

   To shard the database collection properly, see **How Do I Improve Database Performance by Configuring Sharding?**

2. As service data increases, the original database storage is insufficient. You can expand the storage space to fix this problem.

   - To scale up storage for cluster instances, see **Scaling Up a Cluster Instance**.
   - To scale up storage for replica set instances, see **Scaling Up a Replica Set Instance**.
   - To scale up storage for single node instances, see **Scaling Up a Single Node Instance**.

   If the storage space has reached the upper limit of your instance class, change the instance class first.

   - To change the cluster instance class, see **Changing a Cluster Instance Class**.
   - To change the replica set instance class, see **Changing a Replica Set Instance Class**.
   - To change the single node instance class, see **Changing a Single Node Instance Class**.

3. If a large number of expired files occupy the storage space, delete the expired files in time. For example, if the entire database is no longer used, run **dropDatabase** to delete it.

4. The background data processing mechanism is faulty.

   Operations such as write, update, and delete (including index insert and delete) are actually converted to write operations in the background. When data of an instance in use is deleted, the disk space is not reclaimed. Such unreclaimed disk space is called disk fragments. When new data is inserted, these fragments are reused without applying for new disk space. Different underlying storage engines (RocksDB and WiredTiger) vary according to specific scenarios.

   After deleting data, RocksDB directly converts the **delete** operation to append write. After a certain amount of redundant data is accumulated, the background compact thread is automatically triggered to merge and aggregate data of multiple versions to release redundant disk space. You are advised to wait for the system to automatically reclaim the disk space. If the disk space usage is high and close to the **read-only** threshold, contact Huawei technical support.

   After deleting data, WiredTiger merges and aggregates data of multiple versions, causing disk space fragments. However, WiredTiger does not return the disk space to the operating system. WiredTiger marks the disk space for subsequent writes of the current collection, the reserved disk space is preferentially used for subsequent writes of the collection. To release the disk space, run the **compact** command. (Note: This command blocks normal services and is disabled by default.)

4.5 High Memory Usage

If the memory usage of a DDS instance reaches 90% and the swap space usage exceeds 5%, the system responds slowly and even out of memory (OOM) may occur.

This section describes how to fix high memory usage of DB instances.

Viewing the Memory Usage

You can view the monitoring metrics (memory usage and swap usage) to learn the memory usage of instances.

For details, see **Viewing DDS Metrics**.

Figure 4-3 Memory and swap usage

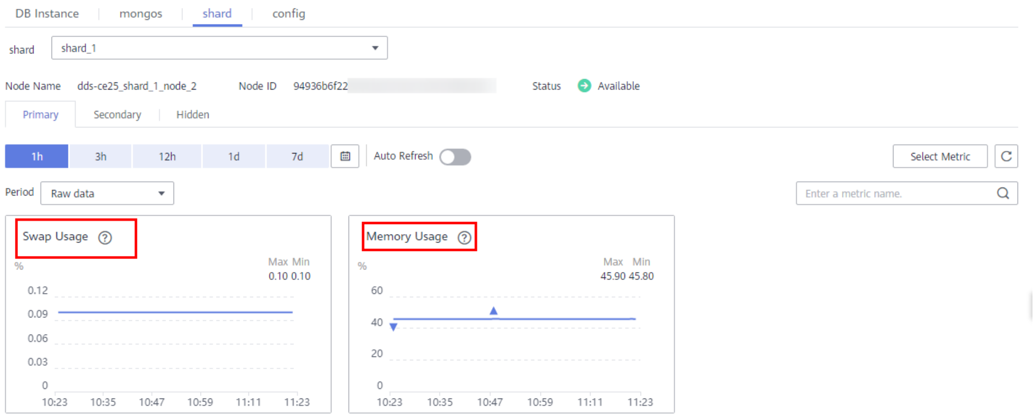

shard: shard_1

Node Name: dds-ce25_shard_1_node_2 Node ID: 94936b6f22[redacted] Status: Available

Period: Raw data

NOTE

By default, 50% memory is reserved, so if the memory usage is 50% but the instance is unloaded, this is normal and you can ignore it.

Solution

1. Control the number of concurrent connections. When connecting to databases, calculate the number of clients and the size of the connection pool configured for each client. The total number of connections cannot exceed 80% of the maximum number of connections supported by the current instance. If there are too many connections, the memory and multi-thread context overhead increases, affecting the delay in request processing.
2. Configure a connection pool. The maximum number of connections in a connection pool is 200.
3. Reduce the memory overhead of a single request. For example, create indexes to reduce collection scanning and memory sorting.
4. If the number of connections remain unchanged but the memory usage keeps increasing, upgrade the memory configuration to prevent system performance deterioration caused by memory overflow and large-scale cache clearing.

- To change cluster instance memory, see **Changing a Cluster Instance Class**.
- To change replica set instance memory, see **Changing a Replica Set Instance Class**.
- To change single node instance memory, see **Changing a Single Node Instance Class**.

4.6 Load Imbalance of Cluster Instances

It is common that load is imbalanced between shard nodes in a cluster instance. If the shard key is incorrectly selected, no chunk is preset, and the load balancing speed between shard nodes is lower than the data insertion speed, load imbalance may occur.

This section describes how to fix load imbalance.

Fault Locating

Step 1 **Connect to a database from the client.**

Step 2 Run the following command to check the shard information:

sh.status()

```
mongos> sh.status()
 \--- Sharding Status ---
   sharding version: {
        "_id" : 1,
        "minCompatibleVersion" : 5,
        "currentVersion" : 6,
        "clusterId" : ObjectId("60f9d67ad4876dd0fe01af84")
   }
   shards:
        {  "_id" : "shard_1",  "host" : "shard_1/172.16.51.249:8637,172.16.63.156:8637",  "state" : 1 }
        {  "_id" : "shard_2",  "host" : "shard_2/172.16.12.98:8637,172.16.53.36:8637",  "state" : 1 }
   active mongoses:
        "4.0.3" : 2
   autosplit:
        Currently enabled: yes
   balancer:
        Currently enabled:  yes
        Currently running:  yes
        Collections with active migrations:
                test.coll started at Wed Jul 28 2021 11:40:41 GMT+0000 (UTC)
        Failed balancer rounds in last 5 attempts:  0
        Migration Results for the last 24 hours:
                300 : Success
   databases:
        {  "_id" : "test",  "primary" : "shard_2",  "partitioned" : true,  "version" : {  "uuid" : UUID("d612d134-
a499-4428-ab21-b53e8f866f67"),  "lastMod" : 1 } }
                test.coll
                        shard key: { "_id" : "hashed" }
                        unique: false
                        balancing: true
                        chunks:
                                shard_1 20
                                shard_2 20
```

- **databases** lists databases for which you enable **enableSharding**.
- **test.coll** is the collection namespace. **test** indicates the name of the database where the collection is located, and **coll** indicates the name of the collection for which sharding is enabled.

- **shard key** is the shard key of the previous collection. **_id**: indicates that the shard is hashed based on **_id**. **_id: -1** indicates that the shard is sharded based on the range of **_id**.
- **chunks** indicates the distribution of shards.

Step 3 Analyze the shard information based on the query result in **Step 2**.

1. If no shard information is queried, the collections are not sharded.

   Run the following command to enable sharding:

   ```
   mongos> sh.enableSharding("<database>")
   mongos> use admin
   mongos> db.runCommand({shardcollection:"<database>.<collection>",key:{"keyname":<value> }})
   ```

2. If an improper shard key is selected, the load may be imbalanced. For example, if a large number of requests are processed on a range of shards, the load on these shards is heavier than other shards, causing load imbalance.

   You can redesign the shard key, for example, changing ranged sharding to hashed sharding.

   ```
   mongos> db.runCommand({shardcollection:"<database>.<collection>",key:{"keyname":<value> }})
   ```

   NOTE

   - If a sharding mode is determined, it cannot be changed easily. The sharding mode must be fully considered in the design phase.
   - For details about how to set data shards, see **How Do I Improve Database Performance by Configuring Sharding?**

3. If a large amount of data is inserted and the data volume exceeds the load capacity of a single shard, shard imbalance occurs and the storage usage of the primary shard is too high.

   You can run the following command to check the network connection of the server and ceck whether the amount of data transmitted by each network adapter reaches the upper limit.

   ```
   sar -n DEV 1  //1 is the interval.
   Average: IFACE  rxpck/s  txpck/s    rxkB/s    txkB/s rxcmp/s txcmp/s rxmcst/s %ifutil
   Average:    lo  1926.94  1926.94  25573.92  25573.92    0.00    0.00     0.00    0.00
   Average:  A1-0     0.00     0.00      0.00      0.00    0.00    0.00     0.00    0.00
   Average:  A1-1     0.00     0.00      0.00      0.00    0.00    0.00     0.00    0.00
   Average:  NIC0     5.17     1.48      0.44      0.92    0.00    0.00     0.00    0.00
   Average:  NIC1     0.00     0.00      0.00      0.00    0.00    0.00     0.00    0.00
   Average:  A0-0  8173.06 92420.66  97102.22 133305.09    0.00    0.00     0.00    0.00
   Average:  A0-1 11431.37  9373.06 156950.45    494.40    0.00    0.00     0.00    0.00
   Average:  B3-0     0.00     0.00      0.00      0.00    0.00    0.00     0.00    0.00
   Average:  B3-1     0.00     0.00      0.00      0.00    0.00    0.00     0.00    0.00
   ```

   NOTE

   - **rxkB/s** is the number of KBs received per second.
   - **txkB/s** is the number of KBs sent per second.

   After the check is complete, press **Ctrl+Z** to exit.

   If the network load is too high, analyze MQL statements, optimize the roadmap, reduce bandwidth consumption, and increase specifications to expand network throughput.

   - Check whether there are sharded collections that do not carry ShardKey. In this case, requests are broadcast, which increases the bandwidth consumption.
   - Control the number of concurrent threads on the client to reduce the network bandwidth traffic.

- If the problem persists, **increase instance specifications** in a timely manner. High-specification nodes can provide higher network throughput.

----End

4.7 Slow Request Locating

In the same service scenario, the query performance depends on the design of the architecture, databases, collections, and indexes. A good design can improve the query performance. On the contrary, a large number of slow queries (statements that take a long time to execute) may occur, which deteriorates system performance.

This document describes the causes and solutions of slow queries.

Fault Locating

DDS allows you to **view slow query logs** on the console. You can start from the slowest operation recorded in the log and optimize the operations one by one.

- If a query takes longer than 1s, the corresponding operation may be abnormal. You need to analyze the problem based on the actual situation.
- If a query takes longer than 10s, the operation needs to be optimized.

NOTE

If an aggregate operation takes more than 10s, it is normal.

Analysis Method

Step 1 Connect to the database.

- To connect to a cluster instance, see **Connecting to a Cluster Instance**.
- To connect to a replica set instance, see **Connecting to a Replica Set Instance**.
- For details about how to connect to a single node instance, see **Connecting to a Single Node Instance**.

Step 2 Run the following command to check the execution plan of a slow query:

explain()

Example:

```
db.test.find({"data_id" : "ae4b5769-896f-465c-9fbd-3fd2f3357637"}).explain();
db.test.find({"data_id" : "775f57c2-b63e-45d7-b581-3822dba231b4"}).explain("executionStats");
```

A covered query does not need to read a document, but directly returns a result from an index, which is very efficient. You can use covering indexes as much as possible. If the output of explain() shows that indexOnly is true, the query is covered by an index.

Step 3 Parse the execution plan.

1. Check the execution time.

   The smaller the values of the following parameters, the better the performance: **executionStats.executionStages.executionTimeMillisEstimate** and **executionStats.executionStages.inputStage.executionTimeMillisEstimate**

Table 4-1 Parameter description

Parameter	Description
executionStats.executionTimeMillis	Execution plan selection and execution time
executionStats.executionStages.executionTimeMillisEstimate	Execution completion time of the execution plan
executionStats.executionStages.inputStage.executionTimeMillisEstimate	Execution completion time of the sub-phase of the execution plan

2. Check the number of scanned records.

   If the three items in **Table 4-2** have the same value, the query performance is the best.

Table 4-2 Parameter description

Parameter	Description
executionStats.nReturned	Number of documents matching the search criteria
executionStats .totalKeysExamined	Number of rows scanned through indexes
executionStats .totalDocsExamined	Number of scanned documents

3. Check the stage status.

   The combinations of stage statuses with better performance are as follows:

   - Fetch+IDHACK
   - Fetch+ixscan,
   - Limit+ (Fetch+ixscan)
   - PROJECTION+ixscan

Table 4-3 Status description

Status Name	Description
COLLSCAN	Full table scan
SORT	In-memory sorting
IDHACK	_id-based query

Status Name	Description
TEXT	Full-text index
COUNTSCAN	Number of unused indexes
FETCH	Index scanning
LIMIT	Using Limit to limit the number of returned records
SUBPLA	$or query stage without using an index
PROJECTION	Number of used indexes
COUNT_SCAN	Number of used indexes

----End

Optimization Plan

- For queries without indexes, create indexes based on the search criteria.
- Hash indexes can be created for point queries.
- Create composite indexes for multi-field queries where a single field is highly repeated.
- Create an ascending or descending index for range lookups with ordered result sets.
- Compound indexes are those indexes sort query results by prefix, so the sequence of query conditions must be the same as that of index fields.
- For partitioned collections (tables) and large collections (with more than 100,000 records), do not use fuzzy query (or do not use LIKE) for tables with a large amount of data. As a result, a large number of records are scanned. You can query data based on the index field, filter out small collections, and then perform fuzzy queries.
- Do not use $not. MongoDB does not index missing data. The $not query requires that all records be scanned in a single result collection. If $not is the only query condition, a full table scan will be performed on the collection.
- If you use $and, put the conditions with the fewest matches before other conditions. If you use $or, put the conditions with the more matches first.
- Check the performance baseline of instance specifications and analyze whether the current service requirements can be met. If the performance bottleneck of the current instance is reached, upgrade the instance specifications in a timely manner.

4.8 Statement Optimization

DDS is inherently a NoSQL database with high performance and strong extensibility. Similar to relational databases, such as RDS for MySQL, RDS for SQL Server, and Oracle, DDS instance performance may also be affected by database design, statement optimization, and index creation.

The following provides suggestions for improving DDS performance in different dimensions:

Creating Databases and Collections

- Use short field names to save storage space. Different from an RDS database, each DDS document has its field names stored in the collection. Short name is recommended.
- Limit the number of documents in a collection to avoid the impact on the query performance. Archive documents periodically if necessary.
- Each document has a default **_id**. Do not change the value of this parameter.
- Capped collections have a faster insertion speed than other collections and can automatically delete old data. You can create capped collections to improve performance based on your service requirements.

Query Operations

Indexes

- Create proper number of indexes for frequently queried fields based on service requirements. Indexes occupy some storage space, and the insert and indexing operations consume resources. It is recommended that the number of indexes in each collection should not exceed 5.
- If data query is slow due to lack of indexes, create proper indexes for frequently queried fields.
- For a query that contains multiple shard keys, create a compound index that contains these keys. The order of shard keys in a compound index is important. A compound index support queries that use the leftmost prefix of the index, and the query is only relevant to the creation sequence of indexes.
- TTL indexes can be used to automatically filter out and delete expired documents. The index for creating TTL must be of type date. TTL indexes are single-field indexes.
- You can create field indexes in a collection. However, if a large number of documents in the collection do not contain key values, you are advised to create sparse indexes.
- When you create text indexes, the field is specified as **text** instead of **1** or **-1**. Each collection has only one text index, but it can index multiple fields.

Command usage

- The findOne method returns the first document that satisfies the specified query criteria from the collection according to the natural order. To return multiple documents, use this method.
- If the query does not require the return of the entire document or is only used to determine whether the key value exists, you can use **$project** to limit the returned field, reducing the network traffic and the memory usage of the client.
- In addition to prefix queries, regular expression queries take longer to execute than using selectors, and indexes are not recommended.
- Some operators that contain **$** in the query may deteriorate the system performance. The following types of operators are not recommended in services. $or, $nin, $not, $ne, and $exists.

Table 4-4 Operator description

Operator	Description
$or	The times of queries depend on the number of conditions. It is used to query all the documents that meet the query conditions in the collection. You are advised to use $in instead.
$nin	Matches most of indexes, and the full table scan is performed.
$not	The query optimizer may fail to match a specific index, and the full table scan is performed.
$ne	Selects the documents where the value of the field is not equal to the specified value. The entire document is scanned.
$exists	Matches each document that contains the field.

For more information, see **official MongoDB documents**.

Precautions

- Indexes cannot be used in operators $where and $exists.
- If the query results need to be sorted, control the number of result sets.
- If multiple field indexes are involved, place the field used for exact match before the index.
- If the key value sequence in the search criteria is different from that in the compound index, DDS automatically changes the query sequence to the same as index sequence.
  - Modification operation

    Modify a document by using operators can improve performance. This method does not need to obtain and modify document data back and forth on the server, and takes less time to serialize and transfer data.
  - Batch insert

    Batch insert can reduce the number of times data is submitted to the server and improve the performance. The BSON size of the data submitted in batches cannot exceed 48 MB.
  - Aggregate operation

    During aggregation, $match must be placed before $group to reduce the number of documents to be processed by the $group operator.

⚠ CAUTION

Improper optimization of slow queries may cause service exceptions.

4.9 Sharding

You can shard a large-size collection for a sharded cluster instance. Sharding distributes data across different machines to make full use of the storage space and compute capability of each shard.

Number of Shards

The following is an example using database **mytable**, collection **mycoll**, and the field **name** as the shard key.

Step 1 Log in to a sharded cluster instance using Mongo Shell.

Step 2 Check whether a collection has been sharded.

```
use <database>
db.<collection>.getShardDistribution()
```

Example:

```
use mytable
db.mycoll.getShardDistribution()
```

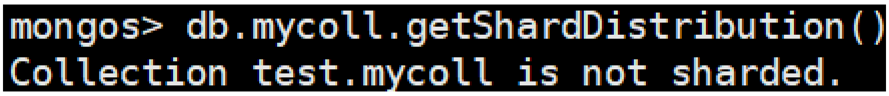

```
mongos> db.mycoll.getShardDistribution()
Collection test.mycoll is not sharded.
```

Step 3 Enable sharding for the databases that belong to the cluster instance.

- Method 1

  ```
  sh.enableSharding("<database>")
  ```

  Example:

  ```
  sh.enableSharding("mytable")
  ```

- Method 2

  ```
  use admin
  db.runCommand({enablesharding:"<database>"})
  ```

Step 4 Shard a collection.

- Method 1

  ```
  sh.shardCollection("<database>.<collection>",{"<keyname>":<value> })
  ```

  Example:

  ```
  sh.shardCollection("mytable.mycoll",{"name":"hashed"},false,{numInitialChunks:5})
  ```

- Method 2

  ```
  use admin
  db.runCommand({shardcollection:"<database>.<collection>",key:{"keyname":<value> }})
  ```

Table 4-5 Parameter description

Parameter	Description
<database>	Database name
<collection>	Collection name.

Parameter	Description
<keyname>	Shard key. Cluster instances are sharded based on the value of this parameter. Select a proper shard key for the collection based on your service requirements. For details, see **Selecting a Shard Key**.
<value>	The sort order based on the range of the shard key. ● 1: Ascending indexes ● -1: Descending indexes ● hashed: indicates that hash sharding is used. Hashed sharding provides more even data distribution across the sharded cluster. For details, see **sh.shardCollection()**.
numInitialChunks	Optional. The minimum number of shards initially created is specified when an empty collection is sharded using a hashed shard key.

Step 5 Check the data storage status of the database on each shard.

```
sh.status()
```

Example:

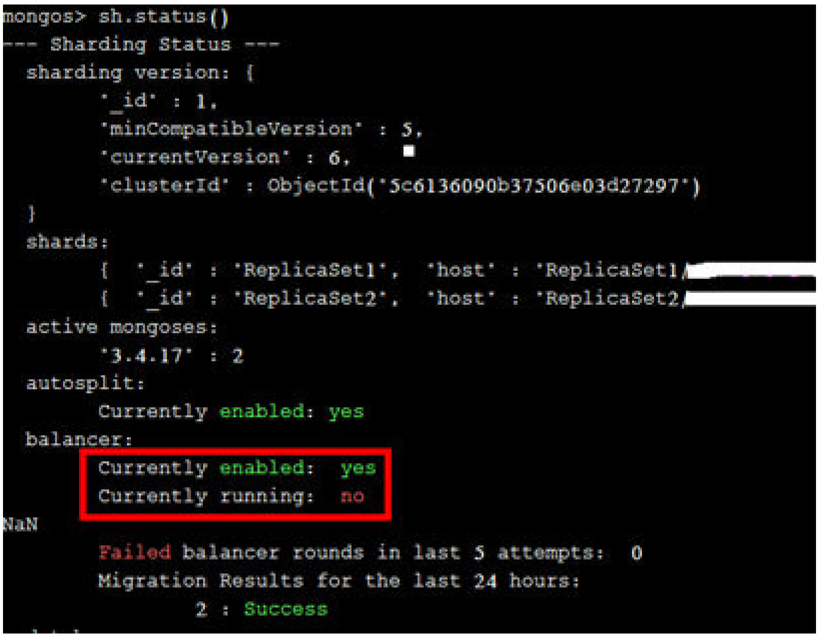

```
mongos> sh.status()
--- Sharding Status ---
  sharding version: {
        "_id" : 1,
        "minCompatibleVersion" : 5,
        "currentVersion" : 6,
        "clusterId" : ObjectId("5c6136090b37506e03d27297")
  }
  shards:
        {  "_id" : "ReplicaSet1",  "host" : "ReplicaSet1/[redacted]
        {  "_id" : "ReplicaSet2",  "host" : "ReplicaSet2/[redacted]
  active mongoses:
        "3.4.17" : 2
  autosplit:
        Currently enabled: yes
  balancer:
        Currently enabled:  yes
        Currently running:  no
NaN
        Failed balancer rounds in last 5 attempts:  0
        Migration Results for the last 24 hours:
                2 : Success
```

----End

Selecting a Shard Key

- **Background**

  Each sharded cluster contains collections as its basic unit. Data in the collection is partitioned by the shard key. Shard key is a field in the collection.

It distributes data evenly across shards. If you do not select a proper shard key, the cluster performance may deteriorate, and the sharding statement execution process may be blocked.

Once the shard key is determined it cannot be changed. If no shard key is suitable for sharding, you need to use a sharding policy and migrate data to a new collection for sharding.

- **Characteristics of proper shard keys**
  - All inserts, updates, and deletes are evenly distributed to all shards in a cluster.
  - The distribution of keys is sufficient.
  - Rare scatter-gather queries.

  If the selected shard key does not have all the preceding features, the read and write scalability of the cluster is affected. For example, If the workload of the find() operation is unevenly distributed in the shards, hot shards will be generated. Similarly, if your write load (inserts, updates, and deletes) is not uniformly distributed across your shards, then you could end up with a hot shard. Therefore, you need to adjust the shard keys based on service requirements, such as read/write status, frequently queried data, and written data.

  After existing data is sharded, if the **filter** filed of the update request does not contain shard keys and **upsert:true** or **multi:false**, the update request will report an error and return message "An upsert on a sharded collection must contain the shard key and have the simple collation.".

- **Judgment criteria**

  You can use the dimensions provided in **Table 4-6** to determine whether the selected shard keys meet your service requirements:

Table 4-6 Reasonable shard keys

Identification Criteria	Description
Cardinality	Cardinality refers to the capability of dividing chunks. For example, if you need to record the student information of a school and use the age as a shard key, data of students of the same age will be stored in only one data segment, which may affect the performance and manageability of your clusters. A much better shard key would be the student number because it is unique. If the student number is used as a shard key, the relatively large cardinality can ensure the even distribution of data.
Write distribution	If a large number of write operations are performed in the same period of time, you want your write load to be evenly distributed over the shards in the cluster. If the data distribution policy is ranged sharding, a monotonically increasing shard key will guarantee that all inserts go into a single shard.

Identification Criteria	Description
Read distribution	Similarly, if a large number of read operations are performed in the same period, you want your read load to be evenly distributed over the shards in a cluster to fully utilize the computing performance of each shard.
Targeted read	The dds mongos query router can perform either a targeted query (query only one shard) or a scatter/gather query (query all of the shards). The only way for the dds mongos to be able to target a single shard is to have the shard key present in the query. Therefore, you need to pick a shard key that will be available for use in the common queries while the application is running. If you pick a synthetic shard key, and your application cannot use it during typical queries, all of your queries will become scatter/gather, thus limiting your ability to scale read load.

Choosing a Distribution Policy

A sharded cluster can store a collection's data on multiple shards. You can distribute data based on the shard keys of documents in the collection.

There are two data distribution policies: ranged sharding and hashed sharding. For details, see **Step 4**.

The following describes the advantages and disadvantages of the two methods.

- **Ranged sharding**

  Ranged-based sharding involves dividing data into contiguous ranges determined by the shard key values. If you assume that a shard key is a line stretched out from positive infinity and negative infinity, each value of the shard key is the mark on the line. You can also assume small and separate segments of a line and that each chunk contains data of a shard key within a certain range.

  Figure 4-4 Distribution of data

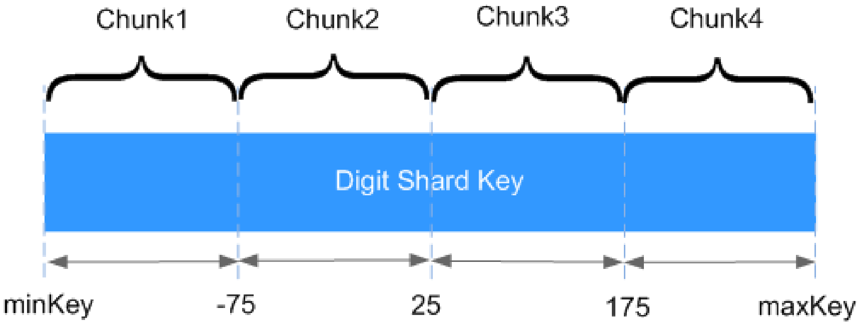

  As shown in the preceding figure, field **x** indicates the shard key of ranged sharding. The value range is [*minKey,maxKey*] and the value is an integer. The value range can be divided into multiple chunks, and each chunk (usually 64 MB) contains a small segment of data. For example, chunk 1 contains all documents in range [minKey, -75] and all data of each chunk is stored on the same shard. That means each shard containing multiple chunks. In addition, the data of each shard is stored on the config server and is evenly distributed by dds mongos based on the workload of each shard.

Ranged sharding can easily meet the requirements of query in a certain range. For example, if you need to query documents whose shard key is in range [-60,20], dds mongos only needs to forward the request to chunk 2.

However, if shard keys are in ascending or descending order, newly inserted documents are likely to be distributed to the same chunk, affecting the expansion of write capability. For example, if **_id** is used as a shard key, the high bits of **_id** automatically generated in the cluster are ascending.

- **Hashed sharding**

  Hashed sharding computes the hash value (64-bit integer) of a single field as the index value; this value is used as your shard key to partition data across your shared cluster. Hashed sharding provides more even data distribution across the sharded cluster because documents with similar shard keys may not be stored in the same chunk.

  Figure 4-5 Distribution of data

  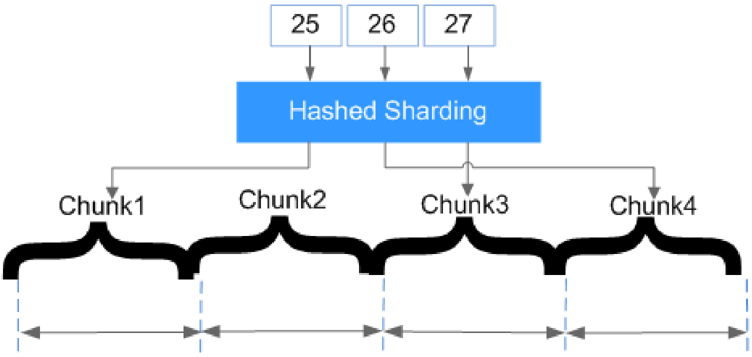

  Hashed sharding randomly distributes documents to each chunk, which fully expands the write capability and makes up for the deficiency of ranged sharding. However, queries in a certain range need to be distributed to all backend shards to obtain documents that meet conditions, resulting in low query efficiency.

5 Permissions Management

5.1 Creating a User and Granting Permissions

This section describes how to use **IAM** to implement fine-grained permissions control for your DDS resources. With IAM, you can:

- Create IAM users for employees based on the organizational structure of your enterprise. Each IAM user has their own security credentials, providing access to DDS resources.
- Grant only the permissions required for users to perform a task.
- Entrust a Huawei account or cloud service to perform professional and efficient O&M on your DDS resources.

If your Huawei account does not need individual IAM users, then you may skip over this topic.

This section describes the procedure for granting permissions (see **Figure 5-1**).

Prerequisites

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

Process Flow

Figure 5-1 Process for granting DDS permissions

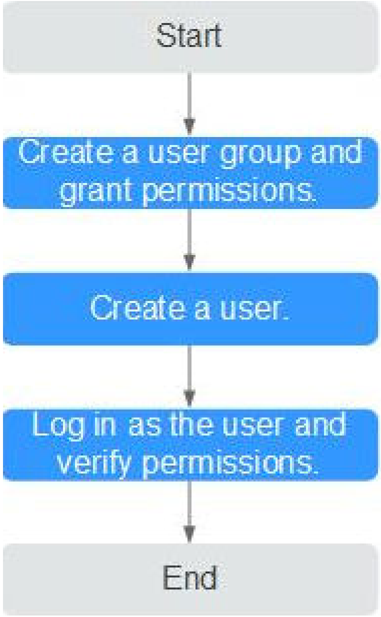
Start
Create a user group and grant permissions.
Create a user.
Log in as the user and verify permissions.
End

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

   Create a user group on the IAM console, and assign the **DDS FullAccess** policy to the group.

   > NOTE
   >
   > To use some interconnected services, you also need to configure permissions of such services.
   >
   > For example, when using DAS to connect to a DB instance, you need to configure the DDS FullAccess and DAS FullAccess permissions.

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

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

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

   Log in to the DDS console by using the newly created user, and verify that the user only has read permissions for DDS.

   Choose **Service List** > **Document Database Service** and click **Buy DB Instance**. If you can buy a DDS DB instance, the required permission policies have taken effect.

5.2 Creating a Custom Policy

Custom policies can be created as a supplement to the system policies of DDS. For the actions supported for custom policies, see **DDS Actions**.

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

- Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.

- JSON: Edit JSON policies from scratch or based on an existing policy.

For details, see **Creating a Custom Policy**. The following section contains examples of common DDS custom policies.

Example Custom Policies

- Example 1: Allowing users to create DDS DB instances

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

- Example 2: Denying DDS DB instance deletion

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

  The following method can be used if you need to assign permissions of the **DDS FullAccess** policy to a user but also forbid the user from deleting DDS DB instances. Create a custom policy for denying DDS DB instance deletion, and assign both policies to the group the user belongs to. Then the user can perform all operations on DDS except deleting DDS instances. The following is an example deny policy:

```
{
    "Version": "1.1",
    "Statement": [
        {
          "Effect": "Deny"
          "Action": [
                "dds:instance:deleteInstance"
            ],
          }
    ]
}
```

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

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

```
{
        "Version": "1.1",
        "Statement": [
                {
                        "Action": [
                                "dds:instance:create",
                                "dds:instance:modify",
                                "dds:instance:deleteInstance",
                 "vpc:publicIps:list",
                 "vpc:publicIps:update"
                        ],
                        "Effect": "Allow"
                }
```

```
    ]
}
Example 4: Setting resource policies
A custom policy can be used to set resource policies, indicating the operation permissions
on the resources under the current action. Currently, the instance name can be configured,
and the asterisk (*) can be used as a wildcard. The following is an example resource policy:
{
    "Version": "1.1",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "dds:instance:list"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "dds:instance:modify"
            ],
            "Resource": [
                "DDS:*:*:instanceName:dds-*"
            ]
        }
    ]
}
```

5.3 Syntax of RBAC Policies

Policy Structure

An RBAC policy consists of a Version, a Statement, and Depends.

Figure 5-2 Policy structure

Version
Statement
Depends

Policy Syntax

Click ⌄ to view the details of a policy. The **DDS Administrator** policy is used as an example to describe the syntax of RBAC policies.

Figure 5-3 DDS Administrator policy

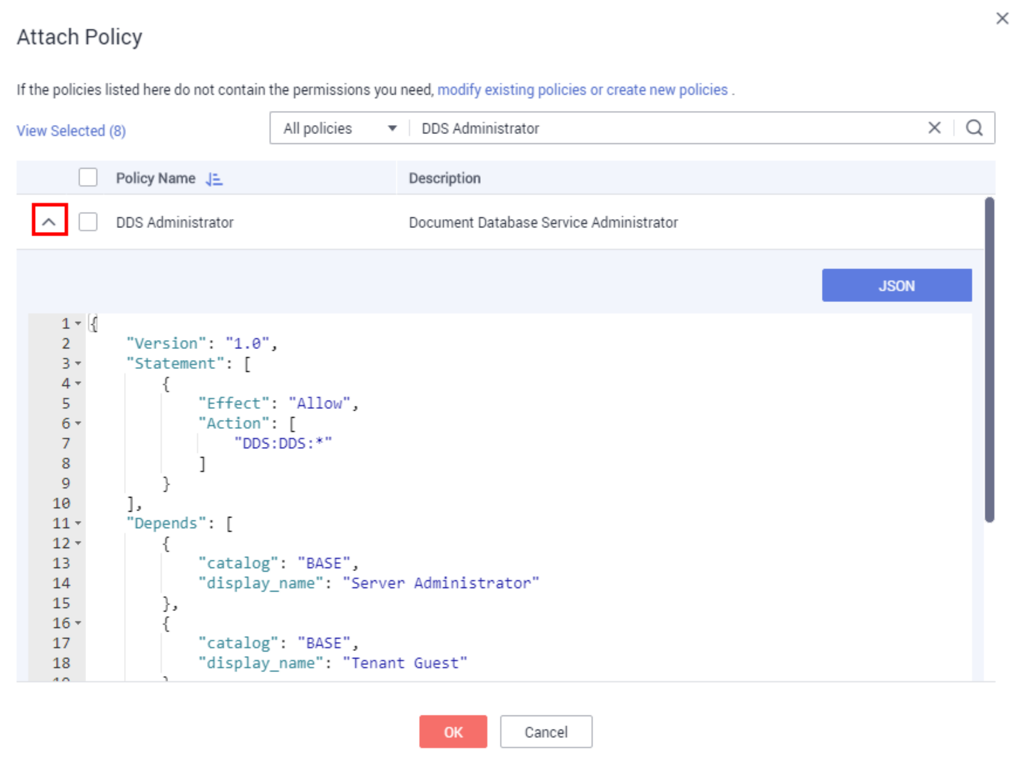

If the policies listed here do not contain the permissions you need, modify existing policies or create new policies .

□	Policy Name	Description
□	DDS Administrator	Document Database Service Administrator

```
{
    "Version": "1.0",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "DDS:DDS:*"
            ]
        }
    ],
    "Depends": [
        {
            "catalog": "BASE",
            "display_name": "Server Administrator"
        },
        {
            "catalog": "BASE",
            "display_name": "Tenant Guest"
```

```
{
        "Version": "1.0",
        "Statement": [
                {
                        "Effect": "Allow",
                        "Action": [
                                "DDS:DDS:*"
                        ],
                        "Resource": [
                                "DDS:*:*:instanceName:dds-*"
                        ],
                }
        ],
        "Depends": [
                {
                        "catalog": "BASE",
                        "display_name": "Server Administrator"
                },
                {
                        "catalog": "BASE",
                        "display_name": "Tenant Guest"
                }
        ]
}
```

Table 5-1 Parameter description

Parameter	Meaning	Value
Version	Policy version	The value is fixed at **1.0**.

Parameter		Meaning	Value
Statement	Action	Operations to be performed on DDS.	Format: *Service name*:*Resource type*:*Operation* **DDS:DDS:***: Permissions for performing all operations on all resource types in DDS.
	Effect	Determines whether the operation defined in an action is allowed.	● Allow ● Deny
	Resource	Defines resource authentication.	This parameter is optional. **DDS:*:*:instanceName:dds-*** indicates that the user has the configured action permissions on all instances whose names start with **dds-**. If this parameter is not specified, the user has the permissions on all instances by default.
Depends	catalog	Name of the service to which dependencies of a policy belong	Service Name Example: **BASE**
	display_name	Name of a dependent policy	Permission name Example: **Server Administrator**

6 Instance Lifecycle Management

6.1 Instance Statuses

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

DB Instance Status

Table 6-1 Status and description

Status	Description
Available	A DB instance is running properly.
Abnormal	A DB instance is faulty.
Creating	A DB instance is being created.
Creation failed	A DB instance fails to be created.
Backing up	An instance backup is being created.
Restarting	A DB instance is being restarted because of a modification that requires restarting it for the modification to take effect.
Switchover in progress	The primary and standby nodes of the replica set instance or the primary and standby shards or configs of a cluster instance are being switched over.
Adding node	shard or dds mongos nodes are being added to a DDS cluster instance.
Deleting node	The node that failed to be added is being deleted.
Scaling up	The storage space of instance nodes is being expanded.
Changing instance class	The CPU or memory of a DB instance is being changed.

Status	Description
Changing to yearly/monthly	The billing mode is being changed from pay-per-use to yearly/monthly.
Checking restoration	The backup of the current DB instance is being restored to a new DB instance.
Restoring	The backup is being restored to the existing DB instance.
Restore failed	Restoring to the existing DB instance failed.
Switching SSL	The SSL channel is being enabled or disabled.
Querying original slow query logs	Show Original Log is being enabled or disabled.
Changing private IP address	The private IP address of a node is being changed.
Changing port	The DB instance port is being changed.
Changing a security group	The security group is being changed.
Frozen	DB instances are frozen when there is no balance in the account.
Minor version upgrade	The minor version upgrade is in progress.
Checking changes	Status of a yearly/monthly instance when the billing mode is being changed.

Parameter Template Status

Table 6-2 Status and description

Status	Description
In-Sync	A database parameter change has taken effect.
Available	Parameters change. Pending restart

6.2 Exporting Instance Information

On the DDS console, you can export information about all DDS instances or information about a specified instance.

NOTE

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Exporting Information of All Instances

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click in the upper right corner of the instance list.

Figure 6-1 Exporting the instance information

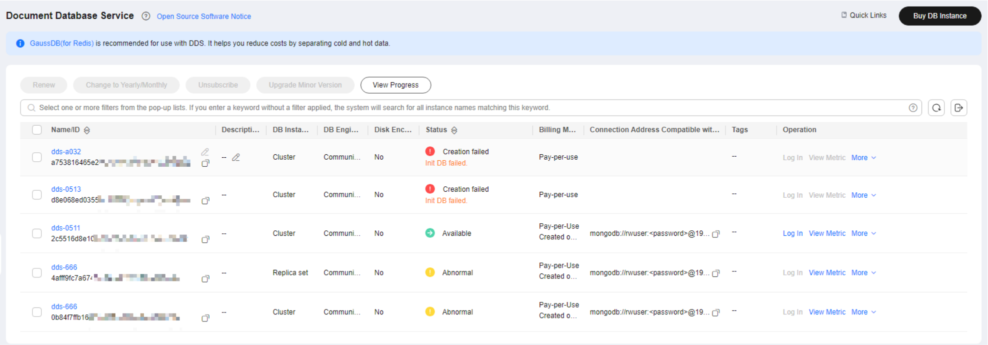

GaussDB(for Redis) is recommended for use with DDS. It helps you reduce costs by separating cold and hot data.

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all instance names matching this keyword.

Name/ID	Descripti...	DB Insta...	DB Engi...	Disk Enc...	Status	Billing M...	Connection Address Compatible wit...	Tags	Operation
dds-a032 a753816465e2[redacted]	--	Cluster	Communi...	No	Creation failed Init DB failed.	Pay-per-use		--	Log In View Metric More
dds-0513 d8e068ed0355[redacted]	--	Cluster	Communi...	No	Creation failed Init DB failed.	Pay-per-use		--	Log In View Metric More
dds-0511 2c5516d8e10[redacted]	--	Cluster	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
dds-666 4afff9fc7a674[redacted]	--	Replica set	Communi...	No	Abnormal	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
dds-666 0b84f7ffb16[redacted]	--	Cluster	Communi...	No	Abnormal	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More

Step 5 In the pop-up box, select the desired items and click **OK**.

Figure 6-2 Export Instance Information

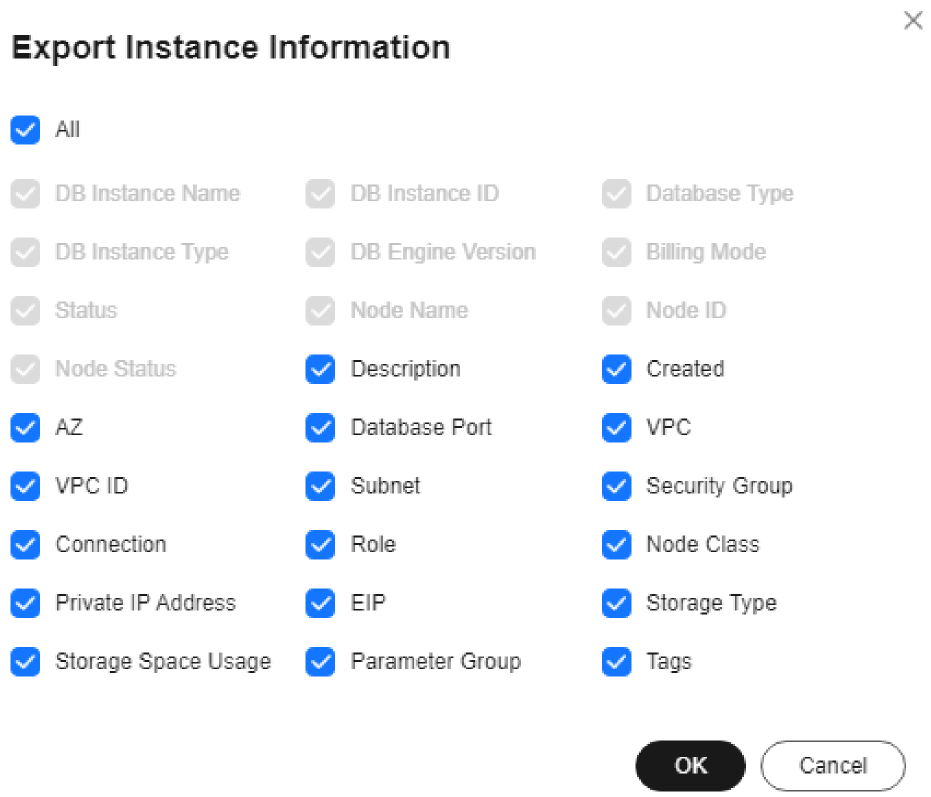

Export Instance Information

☑ All

☑ DB Instance Name	☑ DB Instance ID	☑ Database Type
☑ DB Instance Type	☑ DB Engine Version	☑ Billing Mode
☑ Status	☑ Node Name	☑ Node ID
☑ Node Status	☑ Description	☑ Created
☑ AZ	☑ Database Port	☑ VPC
☑ VPC ID	☑ Subnet	☑ Security Group
☑ Connection	☑ Role	☑ Node Class
☑ Private IP Address	☑ EIP	☑ Storage Type
☑ Storage Space Usage	☑ Parameter Group	☑ Tags

Step 6 View the .xls file exported to your local PC.

----End

Exporting Information of a Specified Instance

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, select the instance and click in the upper right corner of the instance list.

Figure 6-3 Exporting required instance information

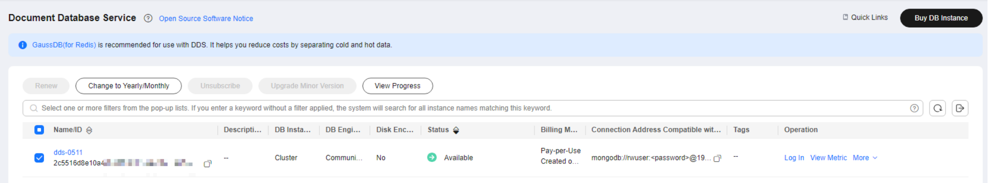

GaussDB(for Redis) is recommended for use with DDS. It helps you reduce costs by separating cold and hot data.

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all instance names matching this keyword.

☐ Name/ID	Descripti...	DB Insta...	DB Engi...	Disk Enc...	Status	Billing M...	Connection Address Compatible wit...	Tags	Operation
☑ dds-0511 2c5516d8e10a4[redacted]	--	Cluster	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More

Step 5 In the pop-up box, select the desired items and click **OK**.

Figure 6-4 Export Instance Information

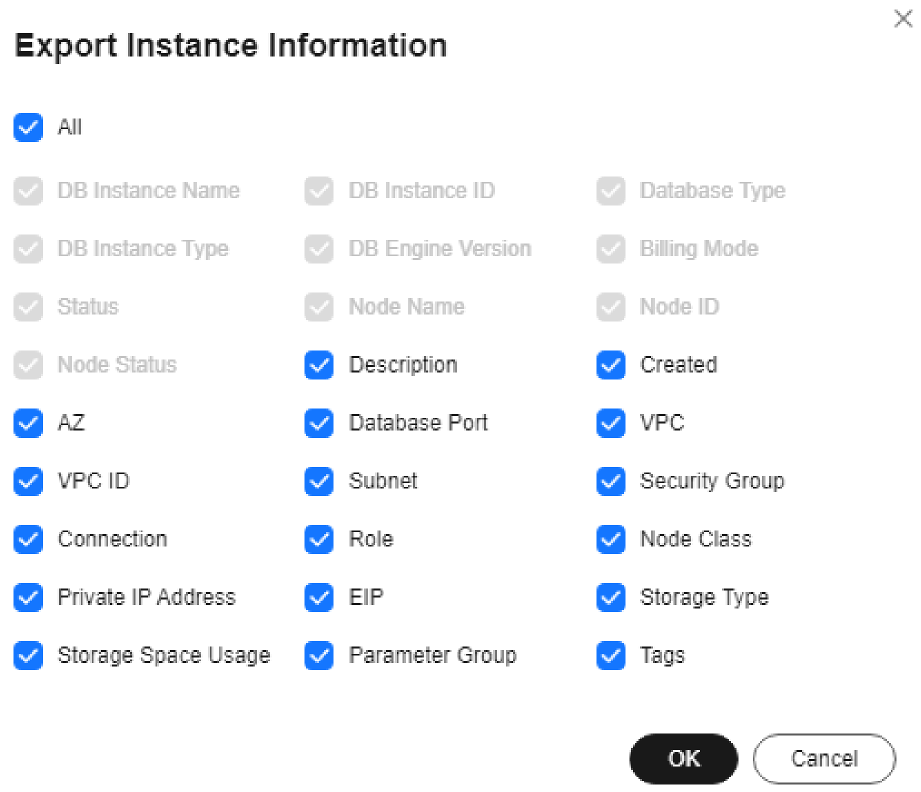

Export Instance Information

☑ All

☑ DB Instance Name ☑ DB Instance ID ☑ Database Type
☑ DB Instance Type ☑ DB Engine Version ☑ Billing Mode
☑ Status ☑ Node Name ☑ Node ID
☑ Node Status ☑ Description ☑ Created
☑ AZ ☑ Database Port ☑ VPC
☑ VPC ID ☑ Subnet ☑ Security Group
☑ Connection ☑ Role ☑ Node Class
☑ Private IP Address ☑ EIP ☑ Storage Type
☑ Storage Space Usage ☑ Parameter Group ☑ Tags

Step 6 View the .xls file exported to your local PC.

----End

6.3 Restarting an Instance or a Node

You may need to occasionally restart an instance to perform routine maintenance. For example, after modifying certain parameters, the instance may need to be restarted to apply the changes.

Precautions

- You can restart an instance only when its status is **Available**.
- Restarting an instance will interrupt services. Exercise caution when performing this operation.
- This instance is not available when it is being restarted. Restarting an instance will clear the cached memory in it. You are advised to restart it during off-peak hours.
- If you restart a cluster or replica set instance, all nodes in the instance are also restarted.
- You can restart a cluster instance or any dds mongos, shard, config node, or read replica in the cluster instance. During the restart, the node cannot be accessed.
- You can restart a replica set instance. During the restart, the instance cannot be accessed.

- You can restart any read replica in a replica set instance. During the restart, the node cannot be accessed.

  NOTE

  Only whitelisted users can restart a read replica in a replica set instance. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

- You can forcibly restart an abnormal node in a DB instance. The node cannot be accessed during the restart.
- After a replica set instance is restarted, the node roles may change.
- It takes less than 30 seconds to start a mongod or dds mongos process. If there are a large number of collections (more than 10,000), it may take several minutes to start the Mongod process. Before the startup is complete, the corresponding node cannot be connected. You are advised to limit the number of collections to less than 10,000 to avoid excessive service loss due to long-time startup.
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see **Operation Protection** in *Identity and Access Management User Guide*.

Restarting an Instance

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the instance and in the **Operation** column, choose **More** > **Restart**.

Figure 6-5 Restarting an instance

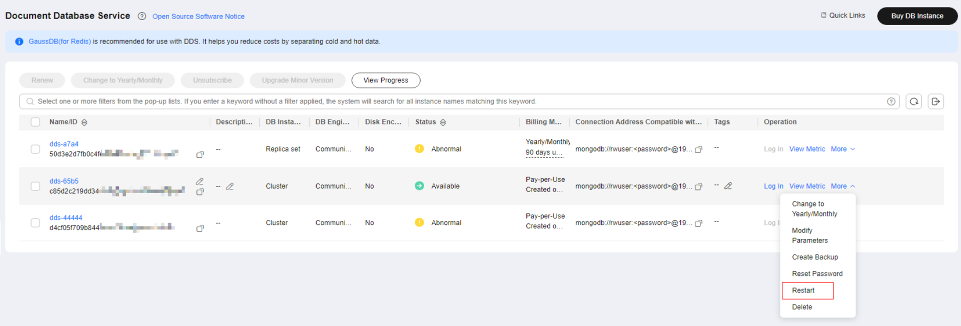

Document Database Service Open Source Software Notice

GaussDB(for Redis) is recommended for use with DDS. It helps you reduce costs by separating cold and hot data.

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all instance names matching this keyword.

Name/ID	Descripti...	DB Insta...	DB Engi...	Disk Enc...	Status	Billing M...	Connection Address Compatible wit...	Tags	Operation
dds-a7a4 50d3e2d7fb0c4f[redacted]	--	Replica set	Communi...	No	Abnormal	Yearly/Monthly 90 days u...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
dds-65b5 c85d2c219dd34[redacted]	--	Cluster	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
dds-44444 d4cf05f709b844[redacted]	--	Cluster	Communi...	No	Abnormal	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log I

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

Figure 6-6 Restarting an instance

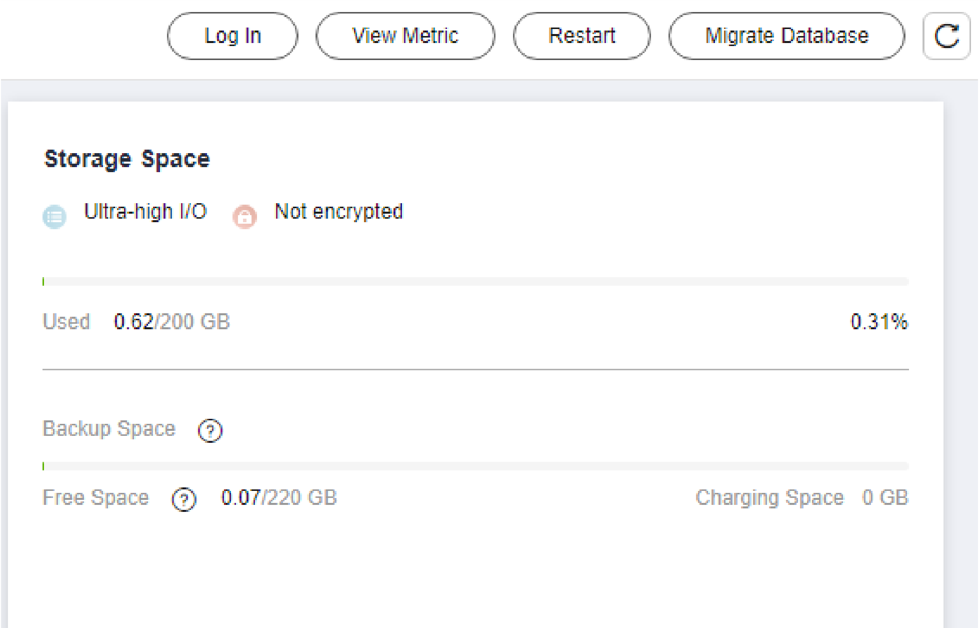

Storage Space

Ultra-high I/O Not encrypted

Used 0.62/200 GB 0.31%

Backup Space

Free Space 0.07/220 GB Charging Space 0 GB

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

Step 6 In the displayed dialog box, click **Yes**.

Step 7 View the instance status.

On the **Instances** page, the instance status is **Restarting**.

----End

Restarting a Cluster Node

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **dds mongos**, **shard**, or **config** tab, locate a node, and in the **Operation** column, click **Restart**.

Figure 6-7 Restarting a dds mongos node

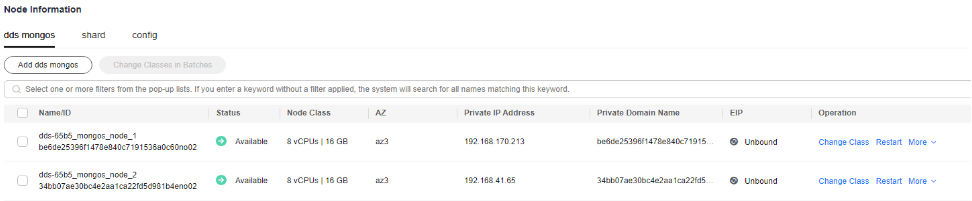

Node Information

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-65b5_mongos_node_1 be6de25396f1478e840c7191536a0c60no02	Available	8 vCPUs | 16 GB	az3	192.168.170.213	be6de25396f1478e840c71915...	Unbound	Change Class Restart More
dds-65b5_mongos_node_2 34bb07ae30bc4e2aa1ca22fd5d981b4eno02	Available	8 vCPUs | 16 GB	az3	192.168.41.65	34bb07ae30bc4e2aa1ca22fd5...	Unbound	Change Class Restart More

Step 6 In the displayed dialog box, click **Yes**.

Step 7 View the node status.

When one node status is **Restarting**, other nodes of the instance cannot be restarted.

----End

Restarting a Read Replica of a Replica Set Instance

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **Read replicas** tab, locate the read replica to be restarted, and click **More** in the **Operation** column.

Figure 6-8 Read replicas

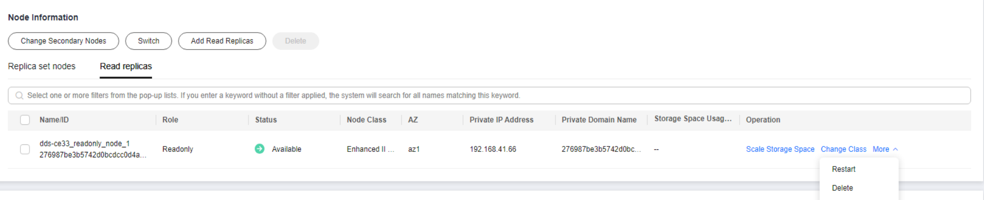

Node Information
Change Secondary Nodes
Switch
Add Read Replicas
Delete
Replica set nodes
Read replicas
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.
Restart
Delete

Name/ID	Role	Status	Node Class	AZ	Private IP Address	Private Domain Name	Storage Space Usag...	Operation
dds-ce33_readonly_node_1 276987be3b5742d0bcdcc0d4a...	Readonly	Available	Enhanced II ...	az1	192.168.41.66	276987be3b5742d0bc...	--	Scale Storage Space Change Class More

Step 6 Select **Restart**.

Step 7 In the displayed dialog box, click **Yes** to restart the read replica.

Step 8 View the status of the read replica.

When one node status is **Restarting**, other nodes of the instance cannot be restarted.

----End

Forcibly Restarting an Abnormal Node

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click its name.

Step 5 In the **Node Information** area on the **Basic Information** page, click **Forcibly Restart** in the **Operation** column of the target abnormal node.

Figure 6-9 Selecting an abnormal node

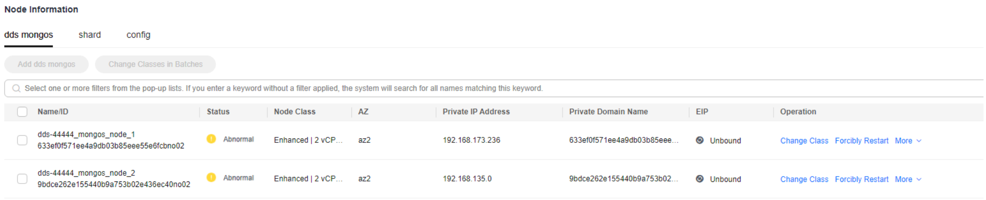

Node Information

dds mongos | shard | config

Add dds mongos | Change Classes in Batches

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

□	Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name	EIP	Operation
□	dds-44444_mongos_node_1 633ef0f571ee4a9db03b85eee55e6fcbno02	Abnormal	Enhanced | 2 vCP...	az2	192.168.173.236	633ef0f571ee4a9db03b85eee...	Unbound	Change Class Forcibly Restart More
□	dds-44444_mongos_node_2 9bdce262e155440b9a753b02e436ec40no02	Abnormal	Enhanced | 2 vCP...	az2	192.168.135.0	9bdce262e155440b9a753b02...	Unbound	Change Class Forcibly Restart More

Step 6 In the displayed dialog box, click **Yes** to restart the abnormal node.

Figure 6-10 Restarting the abnormal node

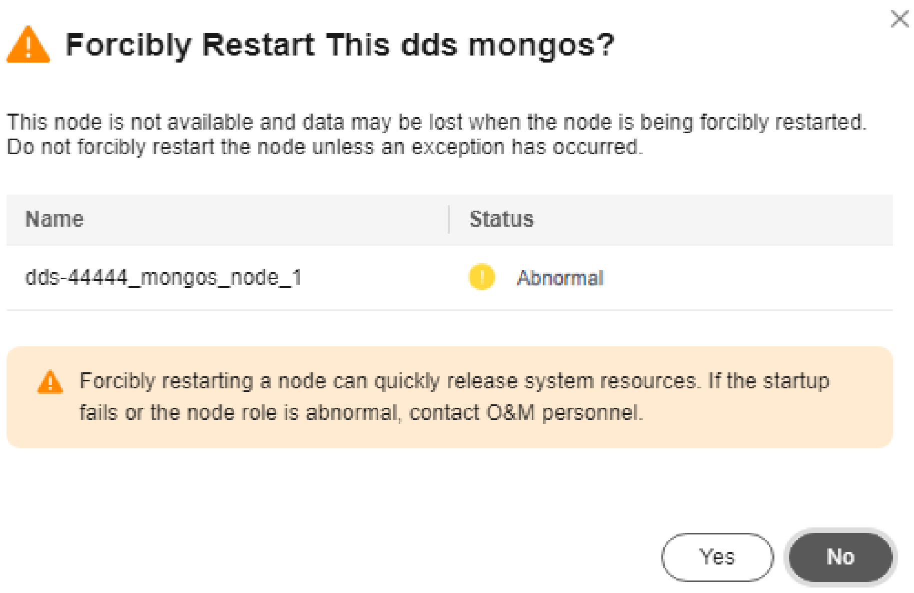

Forcibly Restart This dds mongos?

This node is not available and data may be lost when the node is being forcibly restarted. Do not forcibly restart the node unless an exception has occurred.

Name	Status
dds-44444_mongos_node_1	Abnormal

Forcibly restarting a node can quickly release system resources. If the startup fails or the node role is abnormal, contact O&M personnel.

Step 7 View the status of the node.

When one node status is **Restarting**, other nodes of the instance cannot be restarted.

----End

Restarting Nodes in a Replica Set Instance One by One

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the replica set instance and in the **Operation** column, choose **More** > **Restart**.

Figure 6-11 Restarting a replica set instance

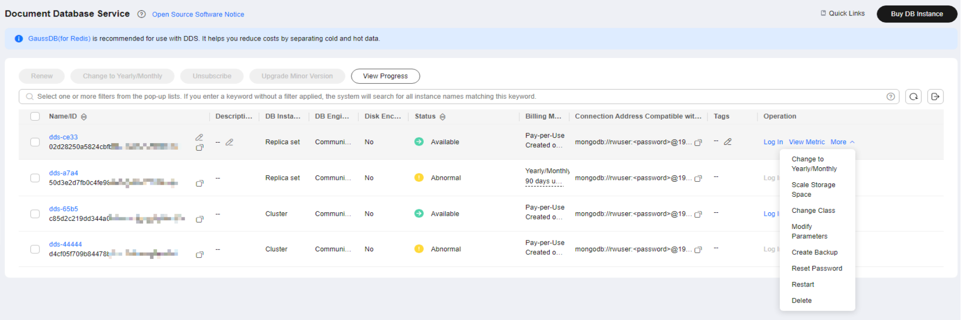

Document Database Service ⓘ Open Source Software Notice

Quick Links | Buy DB Instance

GaussDB(for Redis) is recommended for use with DDS. It helps you reduce costs by separating cold and hot data.

Renew | Change to Yearly/Monthly | Unsubscribe | Upgrade Minor Version | View Progress

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all instance names matching this keyword.

□	Name/ID	Descripti...	DB Insta...	DB Engi...	Disk Enc...	Status	Billing M...	Connection Address Compatible wit...	Tags	Operation
□	dds-ce33 02d28250a5824cbf[redacted]	--	Replica set	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
□	dds-a7a4 50d3e2d7fb0c4fe9[redacted]	--	Replica set	Communi...	No	Abnormal	Yearly/Monthly 90 days u...	mongodb://rwuser:<password>@19...	--	Log I
□	dds-65b5 c85d2c219dd344a[redacted]	--	Cluster	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log I
□	dds-44444 d4cf05f709b84478[redacted]	--	Cluster	Communi...	No	Abnormal	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log I

Alternatively, click the replica set instance name and on the displayed **Basic Information** page, click **Restart** in the upper right corner of the page.

Figure 6-12 Restarting a replica set instance

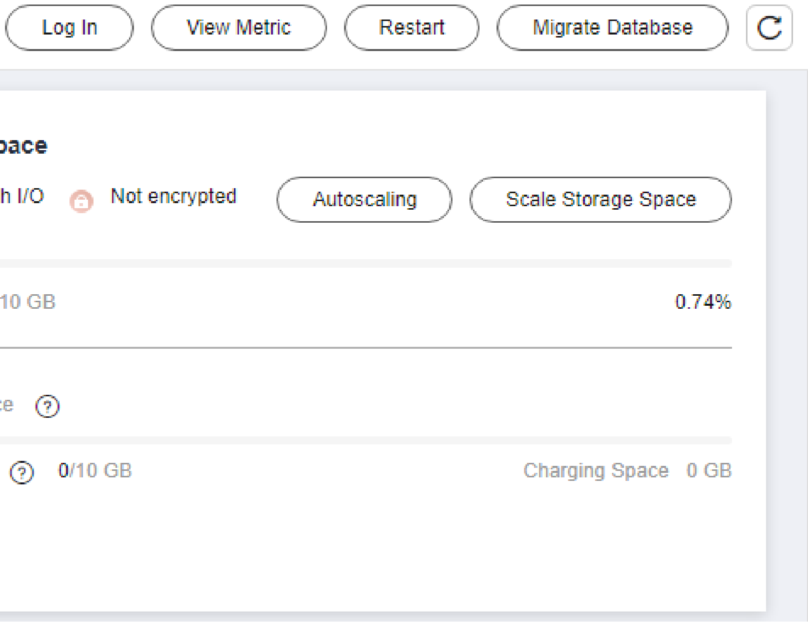

Storage Space

Ultra-high I/O | Not encrypted | Autoscaling | Scale Storage Space

Used 0.07/10 GB 0.74%

Backup Space ?

Free Space ? 0/10 GB Charging Space 0 GB

Step 5 In the displayed dialog box, select **Restart nodes one by one**.

Step 6 Click **Yes** to restart the replica set instance nodes one by one.

Step 7 Check the DB instance status.

On the **Instances** page, the instance status is **Restarting**. If nodes in a replica set instance are restarted one by one, a primary/secondary switchover is triggered.

----End

6.4 Deleting a Pay-per-Use Instance

To delete an instance billed on a pay-per-use basis, you need to locate the instance and click **Delete** on the **Instances** page. After you delete an instance, all of the nodes for that instance are deleted along with it.

Precautions

- To delete an instance billed on a yearly/monthly basis, you need to unsubscribe from the order. For details, see **Billing Termination**.
- After you delete the instance, all its data and all automated backups are automatically deleted as well and cannot be restored. Exercise caution when performing this operation.
- By default, all manual backups are retained in DDS. You can use a backup to restore a deleted instance.
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see **Operation Protection** in *Identity and Access Management User Guide*.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the instance and choose **More** > **Delete** in the **Operation** column.

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

Step 6 In the displayed dialog box, click **Yes**.

----End

6.5 Recycling an Instance

6.5.1 Modifying the Recycling Policy

Unsubscribed yearly/monthly instances and deleted pay-per-use instances can be moved to the recycle bin for management.

Precautions

- The recycling policy is enabled by default and cannot be disabled. Instances in the recycle bin are retained for 7 days by default, and this will not incur any charges.
- Up to 100 instances can be moved to the recycle bin. Once the recycle bin is full, you can still delete instances, but they cannot be placed in the recycle bin, so the deletions will be permanent.
- You can modify the retention period, and the changes only apply to the instances deleted after the changes, so exercise caution when performing this operation.

- Recycling and backup cannot be performed when a node is in the **UNKNOWN** state.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Recycle Bin**.

Step 5 On the **Recycle Bin** page, click **Modify Recycling Policy**. In the displayed dialog box, set the retention period for the deleted DB instances (range: 1 to 7 days). Then, click **OK**.

Figure 6-13 Modify Recycling Policy

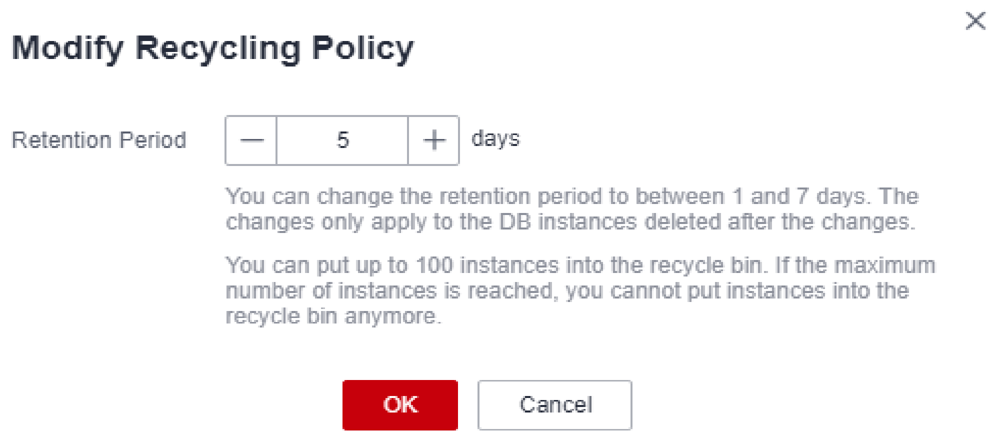

Modify Recycling Policy

Retention Period — 5 + days

You can change the retention period to between 1 and 7 days. The changes only apply to the DB instances deleted after the changes.

You can put up to 100 instances into the recycle bin. If the maximum number of instances is reached, you cannot put instances into the recycle bin anymore.

----End

6.5.2 Rebuilding an Instance

You can rebuild an instance from the recycle bin to restore data.

Precautions

You can rebuild instances from the recycle bin within the retention period.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Recycle Bin**.

Step 5 On the **Recycle Bin** page, locate the instance to be rebuilt and in the **Operation** column, click **Rebuild**.

Figure 6-14 Rebuilding a DB Instance

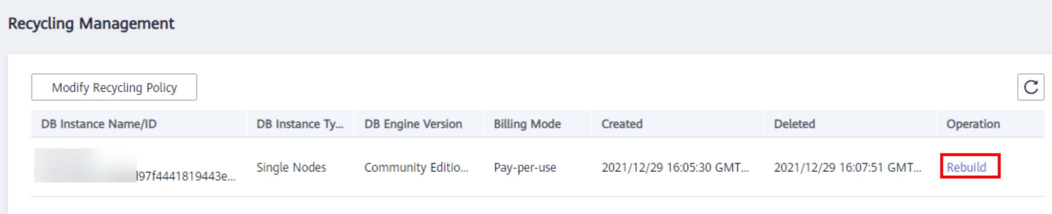

Recycling Management

DB Instance Name/ID	DB Instance Ty...	DB Engine Version	Billing Mode	Created	Deleted	Operation
[redacted] l97f4441819443e...	Single Nodes	Community Editio...	Pay-per-use	2021/12/29 16:05:30 GMT...	2021/12/29 16:07:51 GMT...	Rebuild

Step 6 On the displayed page, set required parameters and submit the rebuilding task. For details, see **Restoring Data to a New Instance**.

----End

7 Instance Modifications

7.1 Changing an Instance Name

This section describes how to change an instance name to identify different instances.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click ✎ next to the instance name you wish to change, enter a new name and click **OK** to apply the changes.

Alternatively, in the **DB Information** area on the **Basic Information** page, click ✎ in the **DB Instance Name** field, enter a new name and click ✓ to apply the changes.

NOTE

- The instance name can be the same as an existing instance name.
- The instance name must contain 4 to 64 characters and must start with a letter. It is case sensitive and can contain letters, digits, hyphens (-), and underscores (_). It cannot contain other special characters.

Step 5 View the results on the **Instances** page.

----End

7.2 Changing an Instance Description

You can add and change descriptions for instances.

> **NOTE**
>
> Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the instance you wish to edit the description for and click ✎ in the **Description** column to edit the instance description. Then, click **OK**.

Alternatively, click the target instance to go to the **Basic Information** page. In the **DB Information** area, click ✎ in the **Description** field to edit the instance description. To submit the change, click ✓.

> **NOTE**
>
> The instance description can contain up to 64 characters, excluding carriage return characters and special characters >!<"&'=

Step 5 View the results on the **Instances** page.

----End

7.3 Modifying an Instance Tag

This section describes how to modify tags of DDS DB instances so that you can filter DB instances by tag.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the instance you wish to edit the tag for and click ✎ in the **Tags** column to edit the instance tag. Then, click **OK**.

Step 5 View the results on the **Instances** page.

----End

7.4 Changing the Name of the Replica Set in the Connection Address

You can change the name of the replica set in the connection address for a DDS DB instance to better meet your service requirements.

Precautions

- This function is available only for replica set instances.
- When you change the replica set name in the connection address of a DDS DB instance, the instance will be unavailable. Exercise caution when performing this operation.
- This operation is not allowed if the DB instance is in any of the following statuses: creating, changing instance class, changing port, restarting, or abnormal.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target replica set instance and click its name.

Step 5 In the **Network Information** area, click ✎ next to **HA Connection Address**.

Figure 7-1 HA Connection Address

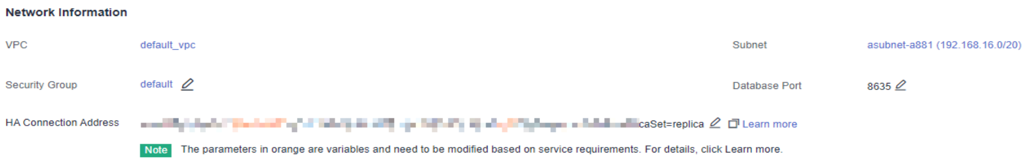

Network Information

VPC default_vpc
Subnet asubnet-a881 (192.168.16.0/20)
Security Group default
Database Port 8635
HA Connection Address [redacted]caSet=replica Learn more
Note The parameters in orange are variables and need to be modified based on service requirements. For details, click Learn more.

Step 6 Enter a new name and click ✓ to save the change.

NOTE

The name of the replica set in the connection address must be 3 to 128 characters long and start with a letter. It is case-sensitive and can contain only letters, digits, and underscores (_).

----End

7.5 Upgrading a Minor Engine Version

DDS supports minor version upgrade to improve performance, add new functions, and fix bugs. For details, see **Kernel Version Description**.

If the database version is a risky version, the system prompts you to upgrade the database patch.

If a new patch is released, you can click **Upgrade Minor Version** on the **Instances** page to upgrade the minor engine version. For details, see **Figure 7-2**.

If the kernel version of your instance has potential risks or major defects, has expired, or has been brought offline, the system will notify you by SMS message or email and deliver an upgrade task during the maintenance window.

Figure 7-2 Minor version upgrade

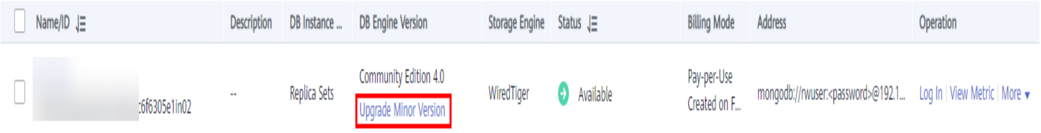

□ Name/ID	Description	DB Instance ...	DB Engine Version	Storage Engine	Status	Billing Mode	Address	Operation
□ [redacted]:6f6305e1in02	--	Replica Sets	Community Edition 4.0 Upgrade Minor Version	WiredTiger	Available	Pay-per-Use Created on F...	mongodb://rwuser:<password>@192.1...	Log In | View Metric | More ▾

Precautions

- A DDS version cannot be downgraded, for example, from 4.0 to 3.4.
- Pay attention to patches that address issues and vulnerabilities from the open source community. When a new patch is released, install the patch in a timely manner.
- During the upgrade, your services may be intermittently interrupted once for up to 30s for each node. Ensure that your instance can be reconnected automatically or perform this operation during off-peak hours.
- DDL operations, such as create event, drop event, and alter event, are not allowed during the upgrade.

Constraints

- Only cluster and replica set instances support minor engine version upgrade.
- This function is available for DB instances of version 3.4 or later.
- If the instance status is abnormal or the instance is being operated, the upgrade cannot be performed.
- The upgrade cannot be performed if the instance nodes are abnormal.

Procedure

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

Step 2 Click ⚲ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the instance you want to upgrade and click **Upgrade Minor Version** in the **DB Engine Version** column.

Figure 7-3 Minor version upgrade

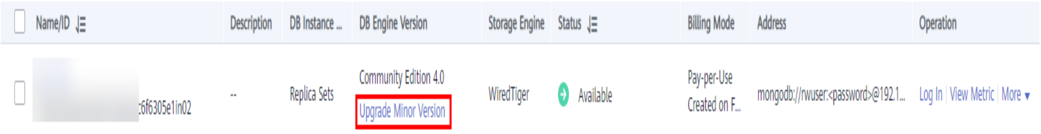

□ Name/ID	Description	DB Instance ...	DB Engine Version	Storage Engine	Status	Billing Mode	Address	Operation
□ [redacted]:6f6305e1in02	--	Replica Sets	Community Edition 4.0 Upgrade Minor Version	WiredTiger	Available	Pay-per-Use Created on F...	mongodb://rwuser:<password>@192.1...	Log In | View Metric | More ▾

Alternatively, click the instance. In the **DB Information** area on the **Basic Information** page, click **Upgrade Minor Version** in the **DB Engine Version** field.

Figure 7-4 Minor version upgrade

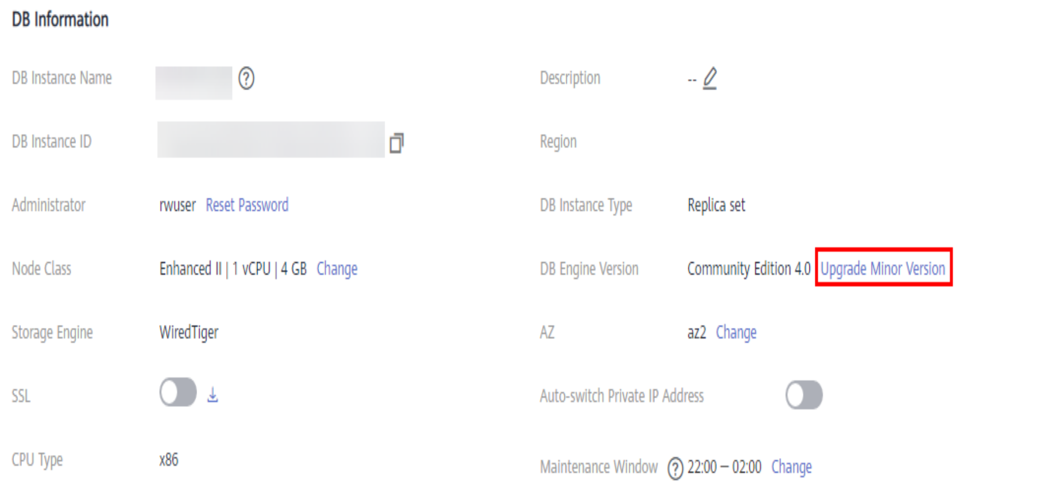

DB Information

DB Instance Name	[redacted]	Description	--
DB Instance ID	[redacted]	Region	
Administrator	rwuser Reset Password	DB Instance Type	Replica set
Node Class	Enhanced II | 1 vCPU | 4 GB Change	DB Engine Version	Community Edition 4.0 Upgrade Minor Version
Storage Engine	WiredTiger	AZ	az2 Change
SSL	○	Auto-switch Private IP Address	○
CPU Type	x86	Maintenance Window	22:00 – 02:00 Change

Step 5 In the displayed dialog box, specify **Scheduled Time** based on service requirements and click **OK**. You can view the upgrade progress on the **Task Center** page.

- **Minimum downtime**: The upgrade has little impact on services.
- **Fastest completion**: The upgrade takes a relatively short time.

Figure 7-5 Selecting a scheduled time

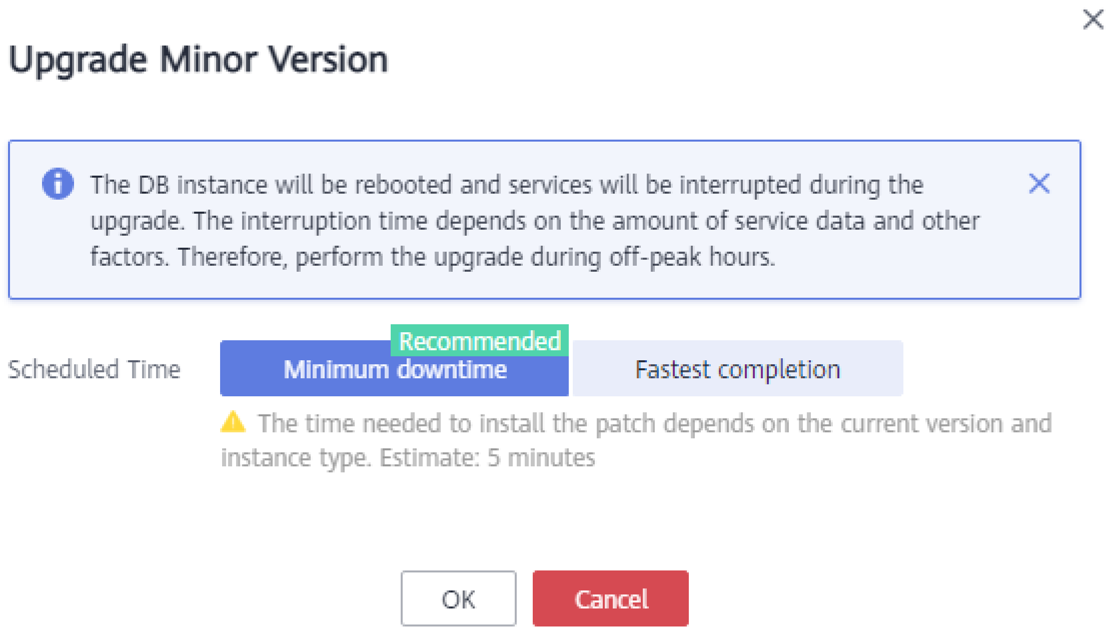

Upgrade Minor Version

The DB instance will be rebooted and services will be interrupted during the upgrade. The interruption time depends on the amount of service data and other factors. Therefore, perform the upgrade during off-peak hours.

Scheduled Time: Minimum downtime (Recommended) | Fastest completion

The time needed to install the patch depends on the current version and instance type. Estimate: 5 minutes

NOTE

The time required for the upgrade varies according to the site requirement.

----End

7.6 Upgrading a Major Engine Version

Precautions

DDS does not support major engine version upgrade on the console. You can use DRS to migrate data as required.

For example, you can use DRS to migrate data from DDS 3.4 to DDS 4.0 without interrupting services.

Constraints

Before migrating data using DRS, you need to create the destination DB instance in advance.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click an instance you want to migrate. On the displayed **Basic Information** page, click **Migrate Database** in the upper right corner of the page.

For more information, see **Migrating Data to the Cloud** in *Data Replication Service User Guide*.

Table 7-1 Database versions

Source DB Version	Destination Database Version	Migration Type
Self-built MongoDB/ Other cloud MongoDB/DDS ● 3.4 ● 4.0 ● 4.2 ● 4.4	DDS ● 3.4 ● 4.0 ● 4.2 ● 4.4	Version upgrade

NOTE

- Data cannot be migrated from a newer version database to an older version database.
- During the specification change, two primary/standby switchovers and two intermittent disconnections will occur. After that, check the DRS task.
- After a major version upgrade, you can change the IP address of the newer version database to the IP address of the older version database. To perform this operation, release the IP address of the older version database first. For details, see **Changing a Private IP Address**.

----End

7.7 Scaling Up Storage Space

7.7.1 Scaling Up a Cluster Instance

You can scale up the storage space of an instance, and the backup space increases accordingly.

- If the purchased storage space exceeds 600 GB and the available storage space is 18 GB, the database will be set to the read-only state when the disk is full.
- If the purchased storage space is less than or equal to 600 GB and the storage usage reaches 97%, the database is set to the read-only state.

In addition, you can set alarm rules for the storage space usage. For details, see **Configuring Alarm Rules**.

For details about the causes and solutions of insufficient storage space, see **High Storage Usage**

Precautions

- Scaling is available when your account balance is sufficient.
- For cluster instances, only shard nodes can be scaled up. dds mongos nodes, config nodes, and read replicas cannot be scaled up.
- If you scale up a DB instance with disks encrypted, the expanded storage space will be encrypted using the original encryption key.
- An instance cannot be scaled up if it is in any of the following statuses:
  - Creating
  - Changing instance class
  - Adding node
  - Deleting node
  - Upgrading minor version
- Services are not interrupted during scaling. The storage type cannot be changed.

Pricing

- A pay-per-use instance is still billed on an hourly basis after the instance is scaled up.
- If you scale up a yearly/monthly instance, you will pay price difference or get a refund.
- For details, see **Product Pricing Details**.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **shard** tab, locate the shard node you want to scale, and click **Scale Storage Space** in the **Operation** column.

Figure 7-6 Scaling up storage space

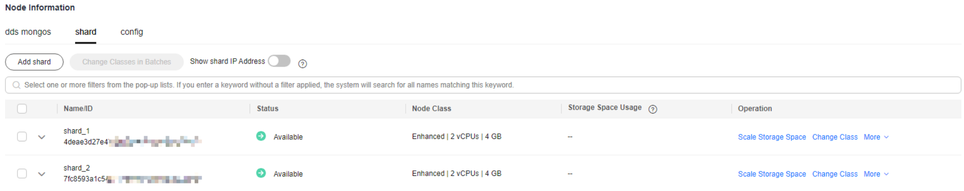
Node Information
dds mongos
shard
config
Add shard
Change Classes in Batches
Show shard IP Address
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

□	Name/ID	Status	Node Class	Storage Space Usage	Operation
□	shard_1 4deae3d27e4[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	--	Scale Storage Space Change Class More
□	shard_2 7fc8593a1c5[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	--	Scale Storage Space Change Class More

Step 6 On the displayed page, specify the desired amount of space to be changed and click **Next**.

Figure 7-7 Scale Storage Space

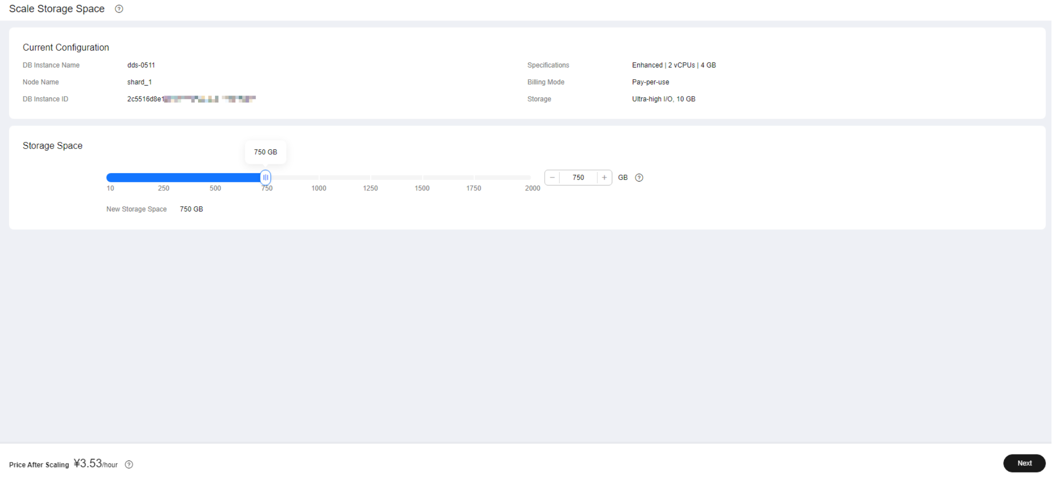

Scale Storage Space

Current Configuration

DB Instance Name: dds-0511

Node Name: shard_1

DB Instance ID: 2c5516d8e1[redacted]

Specifications: Enhanced | 2 vCPUs | 4 GB

Billing Mode: Pay-per-use

Storage: Ultra-high I/O, 10 GB

Storage Space: 750 GB

New Storage Space: 750 GB

Price After Scaling ¥3.53/hour

Select at least 10 GB each time you scale up the storage, and the storage size must be multiples of 10 GB. The maximum amount of storage space is 5,000 GB.

Step 7 On the displayed page, confirm the storage space.

- For yearly/monthly DB instances
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify your settings, click **Submit** to go to the payment page and complete the payment.
- For pay-per-use DB instances
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify the specifications, click **Submit** to scale up the storage space.

Step 8 Check the results.

- This process takes about 3 to 5 minutes. The status of the DB instance in the instance list is **Scaling up**.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.
- In the **Node Information** area on the **Basic Information** page, click the **shard** tab and check whether the scale up was successful.

⚠ CAUTION

If the storage space is scaled up to more than 4 TB, the following risks may occur:

- If there is a large amount of data, the backup task may take a long time or even fail. In this case, the service SLA may be affected. You need to enable snapshot backup to ensure that the backup task can be executed properly. For details about how to enable snapshot backup, see **Enabling or Modifying an Automated Backup Policy**.
- If data is deleted by mistake, it takes a long time to restore a table to a specified point in time or restore a backup to a new instance, affecting the restoration efficiency.
- If the primary/secondary or read-only replication is delayed, it takes a long time to reconnect. As a result, the instance may be disconnected or fail to be reconnected.

----End

Reference

What Should I Do If Storage Usage Is Unusually High?

7.7.2 Scaling Up a Replica Set Instance

You can scale up the storage space of an instance, and the backup space increases accordingly.

- If the purchased storage space exceeds 600 GB and the available storage space is 18 GB, the database will be set to the read-only state when the disk is full.
- If the purchased storage space is less than or equal to 600 GB and the storage usage reaches 97%, the database is set to the read-only state.

In addition, you can set alarm rules for the storage space usage. For details, see **Configuring Alarm Rules**. For details about the causes and solutions of insufficient storage space, see **High Storage Usage**.

Precautions

- Scaling is available when your account balance is sufficient.
- If you scale up a DB instance with disks encrypted, the expanded storage space will be encrypted using the original encryption key.
- An instance cannot be scaled up if it is in any of the following statuses:
  - Creating
  - Changing instance class
  - Adding node
  - Deleting node
  - Upgrading minor version

  - Switchover in progress
- During scaling, services will not be interrupted, and the storage type cannot be changed.

Pricing

- A pay-per-use instance is still billed on an hourly basis after the instance is scaled up.
- If you scale up a yearly/monthly instance, you will pay price difference or get a refund.
- For details, see **Product Pricing Details**.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the replica set instance and choose **More** > **Scale Storage Space** in the **Operation** column.

Figure 7-8 Scale Storage Space

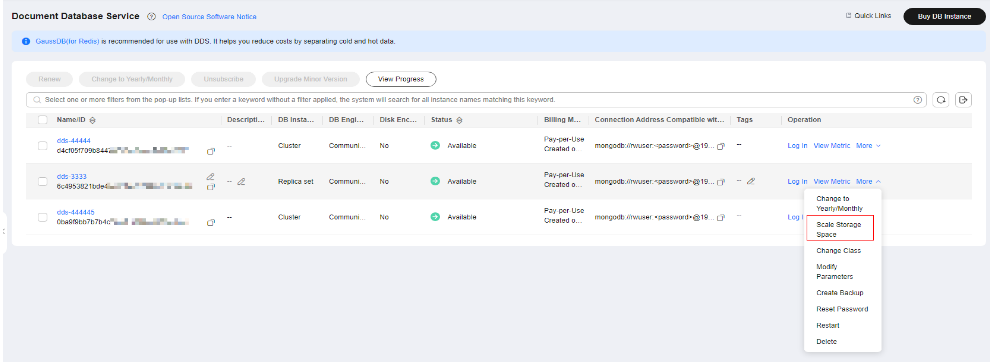

GaussDB(for Redis) is recommended for use with DDS. It helps you reduce costs by separating cold and hot data.

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all instance names matching this keyword.

Name/ID	Descripti...	DB Insta...	DB Engi...	Disk Enc...	Status	Billing M...	Connection Address Compatible wit...	Tags	Operation
dds-44444 d4cf05f709b8447[redacted]	--	Cluster	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
dds-3333 6c4953821bde4[redacted]	--	Replica set	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
dds-444445 0ba9f9bb7b7b4c[redacted]	--	Cluster	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log I

Alternatively, on the **Instances** page, click the name of the replica set instance. In the **Storage Space** area on the **Basic Information** page, click **Scale Storage Space**.

Figure 7-9 Scale Storage Space

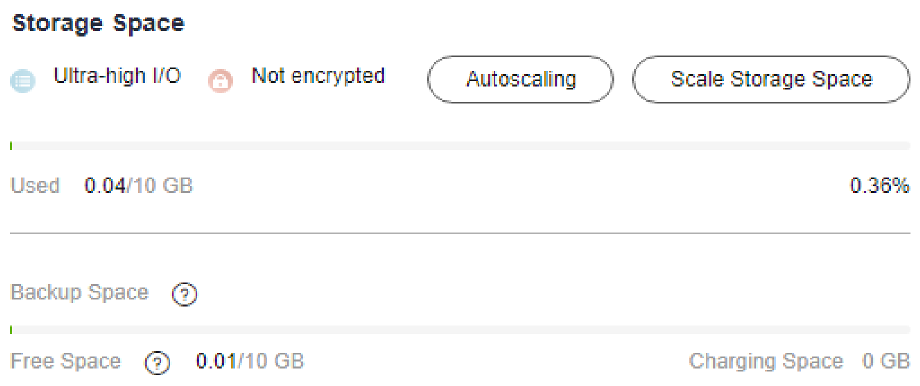

Storage Space

Ultra-high I/O Not encrypted

Used 0.04/10 GB 0.36%

Backup Space

Free Space 0.01/10 GB

Charging Space 0 GB

Step 5 On the displayed page, specify the desired amount of space to be changed and click **Next**.

Figure 7-10 Scale Storage Space

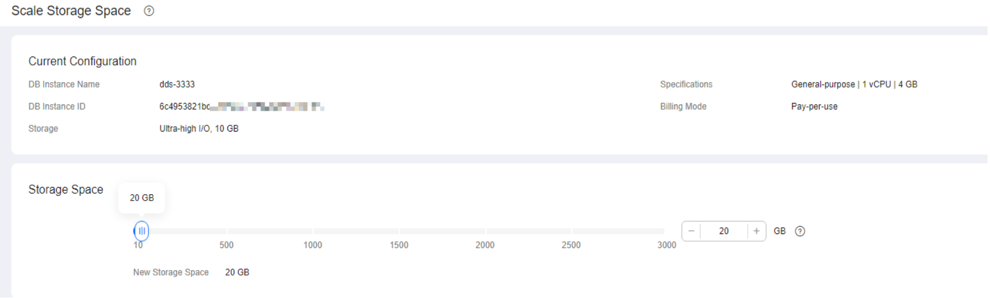

Scale Storage Space

Current Configuration

DB Instance Name dds-3333

DB Instance ID 6c4953821bc[redacted]

Storage Ultra-high I/O, 10 GB

Specifications General-purpose | 1 vCPU | 4 GB

Billing Mode Pay-per-use

Storage Space 20 GB

New Storage Space 20 GB

Select at least 10 GB each time you scale up the storage, and the storage size must be multiples of 10 GB. The maximum amount of storage space is 5,000 GB.

Step 6 On the displayed page, confirm the storage space.

- For yearly/monthly DB instances
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify your settings, click **Submit** to go to the payment page and complete the payment.
- For pay-per-use DB instances
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify the specifications, click **Submit** to scale up the storage space.

Step 7 Check the results.

- This process takes about 3 to 5 minutes. The status of the DB instance in the instance list is **Scaling up**.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.

- In the **Storage Space** area on the **Basic Information** page, check whether the scaling up is successful.

⚠ CAUTION

If the storage space is scaled up to more than 4 TB, the following risks may occur:

- If there is a large amount of data, the backup task may take a long time or even fail. In this case, the service SLA may be affected. You need to enable snapshot backup to ensure that the backup task can be executed properly. For details about how to enable snapshot backup, see **Enabling or Modifying an Automated Backup Policy**.
- If data is deleted by mistake, it takes a long time to restore a table to a specified point in time or restore a backup to a new instance, affecting the restoration efficiency.
- If the primary/secondary or read-only replication is delayed, it takes a long time to reconnect. As a result, the instance may be disconnected or fail to be reconnected.

----End

Reference

What Should I Do If Storage Usage Is Unusually High?

7.7.3 Scaling Up a Read Replica

This section describes how to scale up the storage space of a read replica of a replica set instance.

Precautions

- Scaling is available when your account balance is sufficient.
- If you scale up a DB instance with disks encrypted, the expanded storage space will be encrypted using the original encryption key.
- An instance cannot be scaled up if it is in any of the following statuses:
  - Creating
  - Changing instance class
  - Adding node
  - Deleting node
  - Upgrading minor version
  - Switchover in progress
- During scaling, services will not be interrupted, and the storage type cannot be changed.

Pricing

- A pay-per-use instance is still billed on an hourly basis after the instance is scaled up.

- If you scale up a yearly/monthly instance, you will pay price difference or get a refund.
- For details, see **Product Pricing Details**.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, locate the read replica you want to scale up and click **Scale Storage Space** in the **Operation** column.

Figure 7-11 Scaling storage space

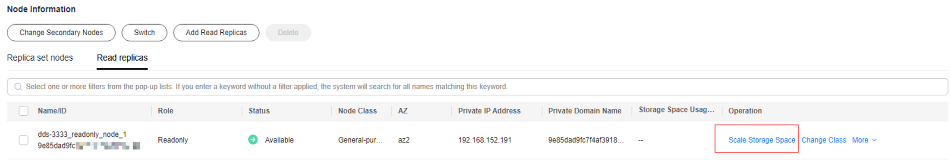

Node Information

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

	Name/ID	Role	Status	Node Class	AZ	Private IP Address	Private Domain Name	Storage Space Usag...	Operation
□	dds-3333_readonly_node_1 9e85dad9fc[redacted]	Readonly	Available	General-pur...	az2	192.168.152.191	9e85dad9fc7f4af3918...	--	Scale Storage Space Change Class More

Step 6 On the displayed page, specify the desired amount of space to be changed and click **Next**.

Figure 7-12 Scaling up a read replica

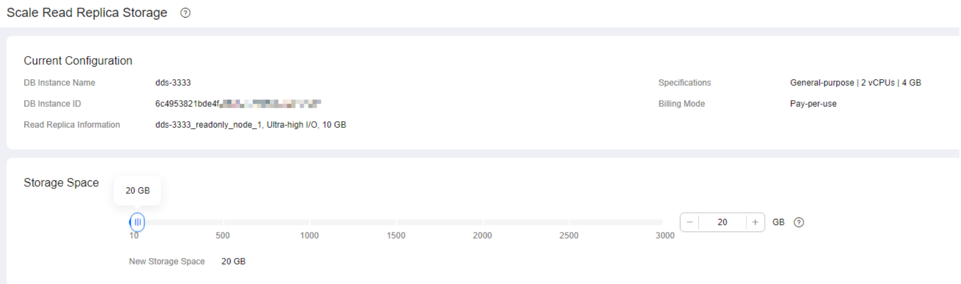

Scale Read Replica Storage

Current Configuration

DB Instance Name	dds-3333	Specifications	General-purpose | 2 vCPUs | 4 GB
DB Instance ID	6c4953821bde4f[redacted]	Billing Mode	Pay-per-use
Read Replica Information	dds-3333_readonly_node_1, Ultra-high I/O, 10 GB		

Storage Space: 20 GB

New Storage Space 20 GB

Select at least 10 GB each time you scale up the storage, and the storage size must be multiples of 10 GB. The maximum amount of storage space is 5,000 GB.

Step 7 On the displayed page, confirm the storage space.

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

  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify the specifications, click **Submit** to scale up the storage space.

Step 8 Check the results.

- This process takes about 3 to 5 minutes. The status of the DB instance in the instance list is **Scaling up**.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.

⚠ CAUTION

If the storage space is scaled up to more than 4 TB, the following risks may occur:

- If there is a large amount of data, the backup task may take a long time or even fail. In this case, the service SLA may be affected. You need to enable snapshot backup to ensure that the backup task can be executed properly. For details about how to enable snapshot backup, see **Enabling or Modifying an Automated Backup Policy**.
- If data is deleted by mistake, it takes a long time to restore a table to a specified point in time or restore a backup to a new instance, affecting the restoration efficiency.
- If the primary/secondary or read-only replication is delayed, it takes a long time to reconnect. As a result, the instance may be disconnected or fail to be reconnected.

----End

Reference

What Should I Do If Storage Usage Is Unusually High?

7.7.4 Scaling Up a Single Node Instance

This section describes how to scale up the storage space of an instance. If you scale up the storage space of an instance, the backup space increases accordingly.

- If the purchased storage space exceeds 600 GB and the available storage space is 18 GB, the database will be set to the read-only state when the disk is full.
- If the purchased storage space is less than or equal to 600 GB and the storage usage reaches 97%, the database is set to the read-only state.

In addition, you can set alarm rules for the storage space usage. For details, see **Configuring Alarm Rules**.

For details about the causes and solutions of insufficient storage space, see **High Storage Usage**

Precautions

- Scaling is available when your account balance is sufficient.
- If you scale up a DB instance with disks encrypted, the expanded storage space will be encrypted using the original encryption key.
- An instance cannot be scaled up if it is in any of the following statuses:
  - Creating
  - Changing instance class
  - Deleting node
  - Upgrading minor version
- Services are not interrupted during scaling. The storage type cannot be changed.

Pricing

- A pay-per-use instance is still billed on an hourly basis after the instance is scaled up.
- If you scale up a yearly/monthly instance, you will pay price difference or get a refund.
- For details, see **Product Pricing Details**.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the single node instance and choose **More** > **Scale Storage Space** in the **Operation** column.

Alternatively, on the **Instances** page, click the name of the single node instance. In the **Storage Space** area on the **Basic Information** page, click **Scale Storage Space**.

Step 5 On the displayed page, specify the desired amount of space to be changed and click **Next**.

Select at least 10 GB each time you scale up the storage, and the storage size must be multiples of 10 GB. The maximum amount of storage space is 1,000 GB.

Step 6 On the displayed page, confirm the storage space.

- For yearly/monthly DB instances
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify your settings, click **Submit** to go to the payment page and complete the payment.
- For pay-per-use DB instances

- If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
- If you do not need to modify the specifications, click **Submit** to scale up the storage space.

Step 7 Check the results.

- This process takes about 3 to 5 minutes. The status of the DB instance in the instance list is **Scaling up**.
- In the upper right corner of the DB instance list, click ↻ to refresh the list. The instance status changes to **Available**.
- In the **Storage Space** area on the **Basic Information** page, check whether the scaling up is successful.

----End

Reference

What Should I Do If Storage Usage Is Unusually High?

7.8 Changing an Instance Class

7.8.1 Changing a Cluster Instance Class

This section describes how to change the class of a cluster instance.

Change Rules

Considering the stability and performance of DDS DB instances, you can change the DB instance class according to the rules listed in **Table 7-2**. Exercise caution when performing this operation.

Table 7-2 Change rules

Original Specification	Target Specification	Supported
General-purpose	General-purpose	√
	Enhanced	×
	Enhanced II	√
Enhanced	General-purpose	√
	Enhanced	×
	Enhanced II	√
Enhanced II	General-purpose	×
	Enhanced	×

Original Specification	Target Specification	Supported
	Enhanced II	√

NOTE

√ indicates that an item is supported, and × indicates that an item is not supported.

Precautions

- An instance cannot be deleted while its instance class is being changed.
- When the instance class is being changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted each time for up to 30s. You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
- After the class of a cluster instance is changed, the system will change the value of **net.maxIncomingConnections** accordingly.
- A maximum of 16 shard nodes can be selected in each batch of class change.
- When the CPU or memory of the shard, config, or dds mongos node in a cluster instance is changed, the read replica class is not changed.
- The classes of read replicas in a cluster instance cannot be changed.
- The classes of yearly/monthly instance shard nodes can only be upgraded or downgraded one at a time.
- Changing the class does not cause data loss.
- If you forcibly change the class of an abnormal node in a DB instance, services may be interrupted.

  NOTE

  To forcibly change the class of an abnormal node, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console.

Pre-check Items for Instance Class Change

- The instance status and the status of the node whose specifications are to be changed are normal.
- The primary/standby replication delay does not exceed 20s. (This pre-check item applies only to shard and config nodes.)

Pricing

- Instances billed on a pay-per-use basis are still billed based on the time used after the instance class is changed.
- If you change the class of a yearly/monthly instance, you will either pay for the difference or receive a refund.

  - If the price of the new instance class is higher than that of the original instance class, you need to pay for the price difference based on the used resource period.
  - If the price of the new instance class is lower than that of the original instance class, you will be refunded the difference based on the used resource period. The refund will be sent to your account. You can click **Billing Center** in the upper right corner of the console to view your account balance.
- For details, see **Product Pricing Details**.

Changing dds mongos Class

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **dds mongos** tab. You can **change the class of a single dds mongos node** or **change the classes of multiple dds mongos nodes at a time**.

- Changing the class of a dds mongos node

  a. In the **Operation** column of the dds mongos node, click **Change Class**.

Figure 7-13 Changing dds mongos class

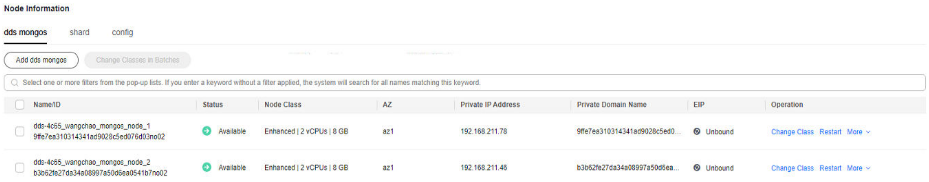

Node Information

dds mongos shard config

Add dds mongos | Change Classes in Batches

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

□	Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name	EIP	Operation
□	dds-4c65_wangchao_mongos_node_1 9ffe7ea310314341ad9028c5ed076d03no02	Available	Enhanced | 2 vCPUs | 8 GB	az1	192.168.211.78	9ffe7ea310314341ad9028c5ed0...	Unbound	Change Class Restart More
□	dds-4c65_wangchao_mongos_node_2 b3b62fe27da34a08997a50d6ea0541b7no02	Available	Enhanced | 2 vCPUs | 8 GB	az1	192.168.211.46	b3b62fe27da34a08997a50d6ea...	Unbound	Change Class Restart More

  b. On the displayed page, select the required specifications, new class, change mode, and scheduled time, and click **Next**.

Figure 7-14 Changing dds mongos class

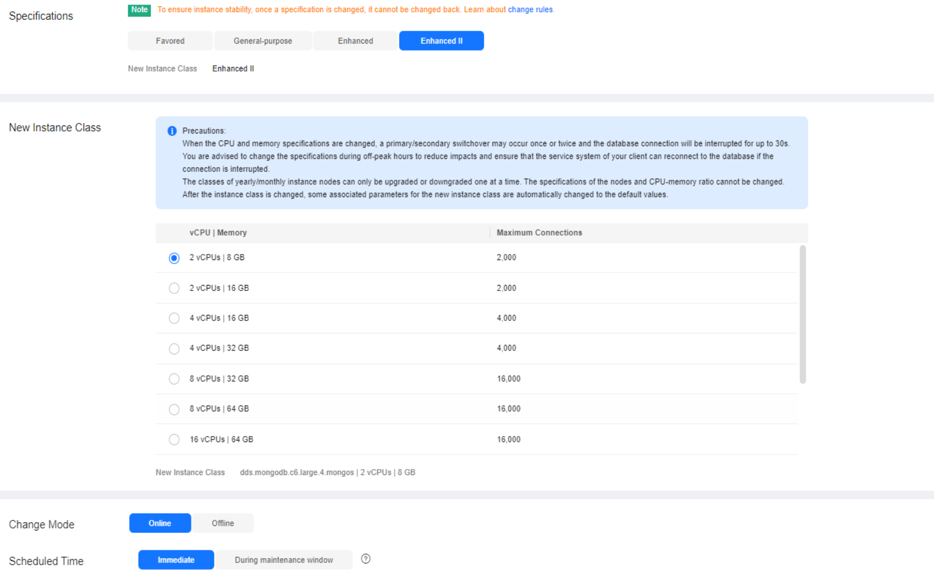

Specifications

Note To ensure instance stability, once a specification is changed, it cannot be changed back. Learn about change rules.

New Instance Class Enhanced II

New Instance Class

Precautions:
When the CPU and memory specifications are changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted for up to 30s. You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
The classes of yearly/monthly instance nodes can only be upgraded or downgraded one at a time. The specifications of the nodes and CPU-memory ratio cannot be changed.
After the instance class is changed, some associated parameters for the new instance class are automatically changed to the default values.

vCPU | Memory	Maximum Connections
● 2 vCPUs | 8 GB	2,000
○ 2 vCPUs | 16 GB	2,000
○ 4 vCPUs | 16 GB	4,000
○ 4 vCPUs | 32 GB	4,000
○ 8 vCPUs | 32 GB	16,000
○ 8 vCPUs | 64 GB	16,000
○ 16 vCPUs | 64 GB	16,000

New Instance Class dds.mongodb.c6.large.4.mongos | 2 vCPUs | 8 GB

Change Mode: [✓ Online] Offline

Scheduled Time: [✓ Immediate] During maintenance window

- Changing the classes of multiple dds mongos nodes in batches
  a. Select the target dds mongos nodes and click **Change Classes in Batches**.

  Figure 7-15 Changing the classes of multiple dds mongos nodes in batches

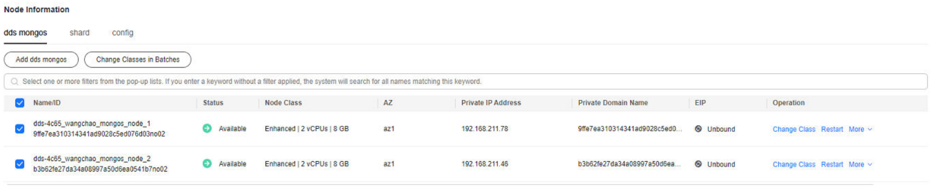

Node Information

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

☑	Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name	EIP	Operation
☑	dds-4c65_wangchao_mongos_node_1 9ffe7ea310314341ad9028c5ed076d03no02	Available	Enhanced | 2 vCPUs | 8 GB	az1	192.168.211.78	9ffe7ea310314341ad9028c5ed0...	Unbound	Change Class Restart More
☑	dds-4c65_wangchao_mongos_node_2 b3b62fe27da34a08997a50d6ea0541b7no02	Available	Enhanced | 2 vCPUs | 8 GB	az1	192.168.211.46	b3b62fe27da34a08997a50d6ea...	Unbound	Change Class Restart More

  b. On the displayed page, select the required specifications, new class, change mode, and scheduled time, and click **Next**.

Figure 7-16 Changing the classes of multiple dds mongos nodes in batches

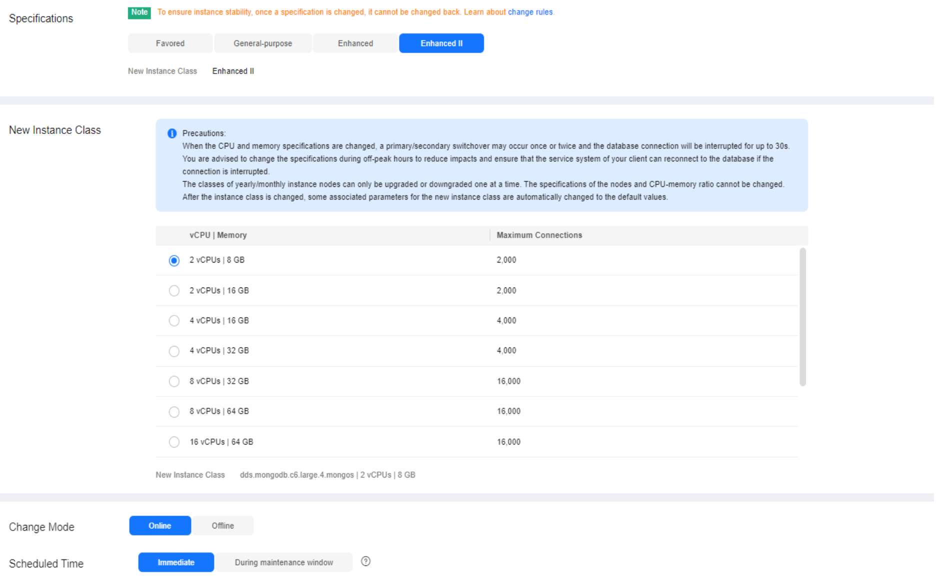

Specifications

Note To ensure instance stability, once a specification is changed, it cannot be changed back. Learn about change rules.

New Instance Class Enhanced II

New Instance Class

Precautions:
When the CPU and memory specifications are changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted for up to 30s. You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
The classes of yearly/monthly instance nodes can only be upgraded or downgraded one at a time. The specifications of the nodes and CPU-memory ratio cannot be changed.
After the instance class is changed, some associated parameters for the new instance class are automatically changed to the default values.

	vCPU | Memory	Maximum Connections
●	2 vCPUs | 8 GB	2,000
○	2 vCPUs | 16 GB	2,000
○	4 vCPUs | 16 GB	4,000
○	4 vCPUs | 32 GB	4,000
○	8 vCPUs | 32 GB	16,000
○	8 vCPUs | 64 GB	16,000
○	16 vCPUs | 64 GB	16,000

New Instance Class dds.mongodb.c6.large.4.mongos | 2 vCPUs | 8 GB

Change Mode: Online (selected) | Offline

Scheduled Time: Immediate (selected) | During maintenance window

NOTE

- Online change: The specifications of multiple dds mongos nodes will be changed one by one. The time required depends on the number of instance nodes whose specifications need to be changed. Each node takes about 5 to 10 minutes. You are advised to connect to a DB instance using the HA connection address and ensure that your applications support automatic reconnection.
- Offline change: The specifications of multiple dds mongos nodes will be changed concurrently and the database will be unavailable during the specification change. It takes about 5 to 10 minutes.
- The specifications change of dds mongos nodes does not involve primary/standby switchover.

Step 6 On the displayed page, confirm the class.

- If you need to modify your settings, click **Previous**.
- For pay-per-use instances

  If you do not need to modify your settings, click **Submit** to change the class. After the specifications are changed, you are still charged on an hourly basis.
- For yearly/monthly instances
  - If you intend to scale down the class, click **Submit**. The refund is automatically returned to your account.
  - If you intend to scale up the class, click **Pay Now**. The scaling starts only after the payment is successful.

Step 7 View the results.

- When the instance class is being changed, the status displayed in the **Status** column is **Changing instance class**. This process takes about 10 minutes.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.

- In the **Node Information** area on the **Basic Information** page, click the **dds mongos** tab and view the new class.

----End

Changing shard Class

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **shard** tab. You can change the class of a single shard or change the classes of multiple shards at a time.

- Changing the class of a shard
  a. In the **Operation** column of the shard node, click **Change Class**.

Figure 7-17 Changing the class of a shard

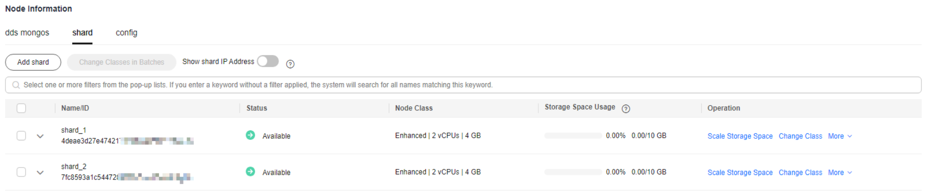

Node Information

dds mongos shard config

Add shard | Change Classes in Batches | Show shard IP Address

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Status	Node Class	Storage Space Usage	Operation
shard_1 4deae3d27e47421[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	0.00% 0.00/10 GB	Scale Storage Space Change Class More
shard_2 7fc8593a1c54472[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	0.00% 0.00/10 GB	Scale Storage Space Change Class More

  b. On the displayed page, select the required specifications and new class and click **Next**.

NOTE

The time required depends on the number of instance nodes whose specifications are to be changed. It takes about 5 to 10 minutes for each node. When the instance class is being changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted each time for up to 30s. Before the specification change, learn about Pre-check Items for Instance Class Change. You are advised to change the class during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.

Figure 7-18 Changing the class of a shard

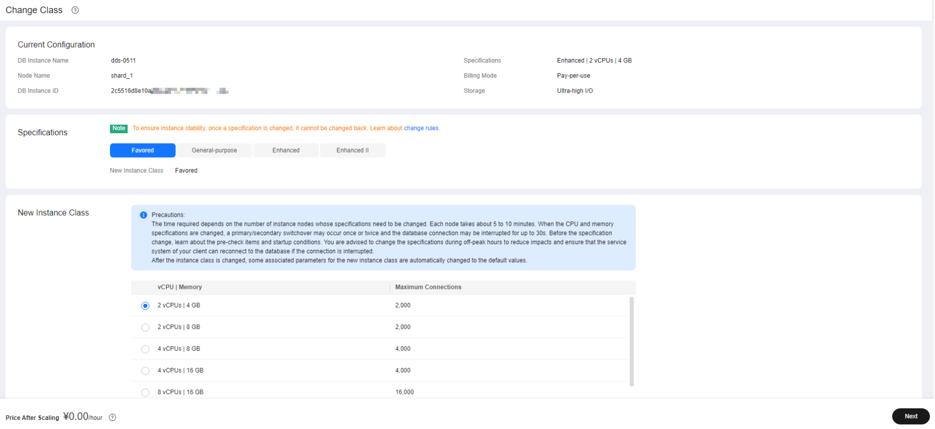

Change Class

Current Configuration

DB Instance Name	dds-0511	Specifications	Enhanced | 2 vCPUs | 4 GB
Node Name	shard_1	Billing Mode	Pay-per-use
DB Instance ID	2c5516d8e10a[redacted]	Storage	Ultra-high I/O

Specifications

Note To ensure instance stability, once a specification is changed, it cannot be changed back. Learn about change rules.

New Instance Class Favored

New Instance Class

Precautions:
The time required depends on the number of instance nodes whose specifications need to be changed. Each node takes about 5 to 10 minutes. When the CPU and memory specifications are changed, a primary/secondary switchover may occur once or twice and the database connection may be interrupted for up to 30s. Before the specification change, learn about the pre-check items and startup conditions. You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
After the instance class is changed, some associated parameters for the new instance class are automatically changed to the default values.

	vCPU | Memory	Maximum Connections
●	2 vCPUs | 4 GB	2,000
○	2 vCPUs | 8 GB	2,000
○	4 vCPUs | 8 GB	4,000
○	4 vCPUs | 16 GB	4,000
○	8 vCPUs | 16 GB	16,000

Price After Scaling ¥0.00/hour

- Changing the classes of multiple shards in batches
  a. Select the target shards and click **Change Classes in Batches**.

Figure 7-19 Changing the classes of multiple shards in batches

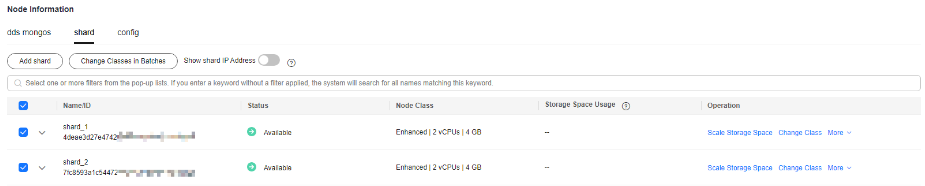

Node Information

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

☑	Name/ID	Status	Node Class	Storage Space Usage	Operation
☑	shard_1 4deae3d27e4742[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	--	Scale Storage Space Change Class More
☑	shard_2 7fc8593a1c54472[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	--	Scale Storage Space Change Class More

  b. On the displayed page, select the required specifications and new class and click **Next**.

NOTE

The time required depends on the number of instance nodes whose class is to be changed. It takes about 5 to 10 minutes for each node. When the instance class is being changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted each time for up to 30s. Before the specification change, learn about **Pre-check Items for Instance Class Change**. You are advised to change the class during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.

Figure 7-20 Changing the classes of multiple shards in batches

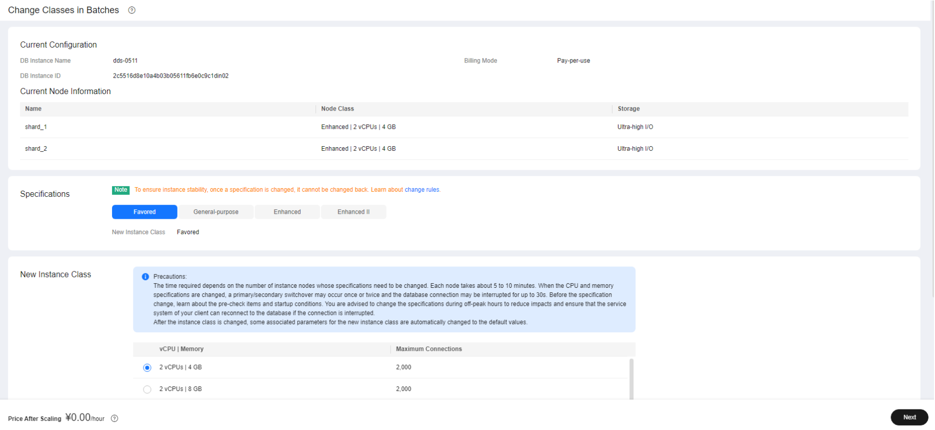

Change Classes in Batches

Current Configuration

DB Instance Name	dds-0511	Billing Mode	Pay-per-use
DB Instance ID	2c5516d8e10a4b03b05611fb6e0c9c1din02		

Current Node Information

Name	Node Class	Storage
shard_1	Enhanced | 2 vCPUs | 4 GB	Ultra-high I/O
shard_2	Enhanced | 2 vCPUs | 4 GB	Ultra-high I/O

Specifications

Note To ensure instance stability, once a specification is changed, it cannot be changed back. Learn about change rules.

New Instance Class Favored

New Instance Class

Precautions:
The time required depends on the number of instance nodes whose specifications need to be changed. Each node takes about 5 to 10 minutes. When the CPU and memory specifications are changed, a primary/secondary switchover may occur once or twice and the database connection may be interrupted for up to 30s. Before the specification change, learn about the pre-check items and startup conditions. You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
After the instance class is changed, some associated parameters for the new instance class are automatically changed to the default values.

	vCPU | Memory	Maximum Connections
●	2 vCPUs | 4 GB	2,000
○	2 vCPUs | 8 GB	2,000

Price After Scaling ¥0.00/hour

Step 6 On the displayed page, confirm the class.

- If you need to modify your settings, click **Previous**.
- For pay-per-use instance

  If you do not need to modify your settings, click **Submit** to change the class. After the specifications are changed, you are still charged on an hourly basis.
- For yearly/monthly instance:
  - If you intend to scale down the class, click **Submit**. The refund is automatically returned to your account.
  - If you intend to scale up the class, click **Pay Now**. The scaling starts only after the payment is successful.

Step 7 View the results.

- When the instance class is being changed, the status displayed in the **Status** column is **Changing instance class**. This process takes about 25 to 30 minutes.

  NOTE

  High database load increases the specification change duration. You are advised to change the specifications during off-peak hours to reduce impacts.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.
- Go to the **Basic Information** page of the cluster instance you scaled up, click the **shard** tab in the **Node Information** area, and view the new class.

----End

Changing config Class

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **config** tab, locate the config node, and click **Change Class** in the **Operation** column.

Figure 7-21 Changing config class

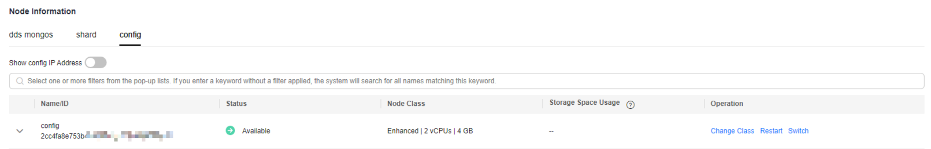

Node Information

dds mongos shard config

Show config IP Address

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Status	Node Class	Storage Space Usage	Operation
config 2cc4fa8e753b[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	--	Change Class Restart Switch

Step 6 On the displayed page, select the required specifications and new class and click **Next**.

📖 NOTE

The time required depends on the number of instance nodes whose specifications are to be changed. It takes about 5 to 10 minutes for each node. When the instance class is being changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted each time for up to 30s. Before the specification change, learn about **Pre-check Items for Instance Class Change**. You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.

Figure 7-22 Changing config class

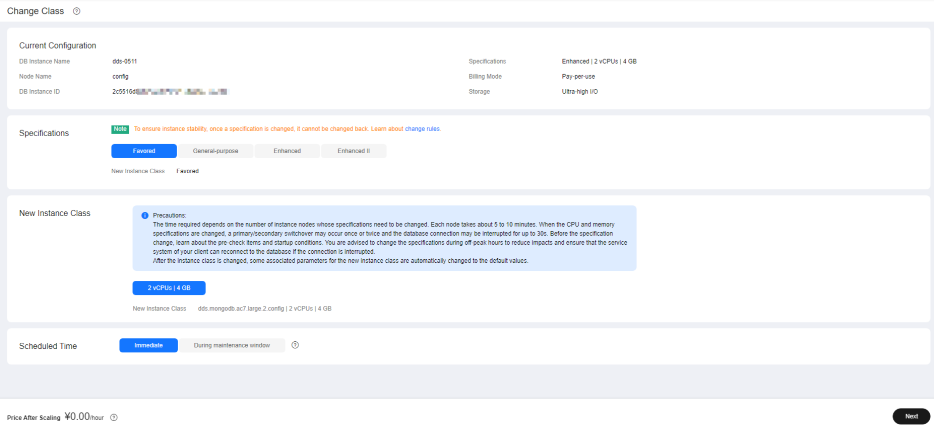

Change Class

Current Configuration

DB Instance Name dds-0511

Node Name config

DB Instance ID 2c5516d[redacted]

Specifications Enhanced | 2 vCPUs | 4 GB

Billing Mode Pay-per-use

Storage Ultra-high I/O

Specifications

Note To ensure instance stability, once a specification is changed, it cannot be changed back. Learn about change rules.

New Instance Class Favored

New Instance Class

Precautions:
The time required depends on the number of instance nodes whose specifications need to be changed. Each node takes about 5 to 10 minutes. When the CPU and memory specifications are changed, a primary/secondary switchover may occur once or twice and the database connection may be interrupted for up to 30s. Before the specification change, learn about the pre-check items and startup conditions. You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
After the instance class is changed, some associated parameters for the new instance class are automatically changed to the default values.

2 vCPUs | 4 GB

New Instance Class dds.mongodb.ac7.large.2.config | 2 vCPUs | 4 GB

Scheduled Time Immediate

Price After Scaling ¥0.00/hour

Step 7 On the displayed page, confirm the class.

- If you need to modify your settings, click **Previous**.
- For pay-per-use instances

  If you do not need to modify your settings, click **Submit** to change the class. After the specifications are changed, you are still charged on an hourly basis.
- For yearly/monthly instances
  - If you intend to scale down the class, click **Submit**. The refund is automatically returned to your account.
  - If you intend to scale up the class, click **Pay Now**. The scaling starts only after the payment is successful.

Step 8 View the results.

- When the instance class is being changed, the status displayed in the **Status** column is **Changing instance class**. This process takes about 25 to 30 minutes.

  📖 NOTE

  High database load increases the specification change duration. You are advised to change the specifications during off-peak hours to reduce impacts.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.
- Go to the **Basic Information** page of the cluster instance you scaled up, click the **config** tab in the **Node Information** area, and view the new class.

----End

Reference

How Do I Solve the High CPU Usage Issue?

7.8.2 Changing a Replica Set Instance Class

This section describes how to change the class of a replica set instance.

Change Rules

Considering the stability and performance of DDS DB instances, you can change the DB instance class according to the rules listed in **Table 7-3**. Exercise caution when performing this operation.

Table 7-3 Change rules

Original Specification	Target Specification	Supported
General-purpose	General-purpose	√
	Enhanced	×
	Enhanced II	√
Enhanced	General-purpose	√
	Enhanced	×
	Enhanced II	√
Enhanced II	General-purpose	×
	Enhanced	×
	Enhanced II	√

NOTE

√ indicates that an item is supported, and × indicates that an item is not supported.

Precautions

- A DB instance cannot be deleted while its instance class is being changed.
- When the CPU or memory of a replica set instance is changed, the read replica class is not changed.
- When the instance class is being changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted each time for up to 30s. You are advised to change the class during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
- After the class of a replica set instance is changed, the system will change the value of **net.maxIncomingConnections** accordingly.

- Changing the class does not cause data loss.
- If you forcibly change the class of an abnormal node in a DB instance, services may be interrupted.

  NOTE

  To forcibly change the class of an abnormal node, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console.

Pre-check Items for Instance Class Change

- The instance status and the status of the node whose class is to be changed are normal.
- The primary/standby replication delay does not exceed 20s.

Billing

- Instances in pay-per-use mode are still charged based on the time used after the instance class is changed.
- If you change the class of a yearly/monthly instance, you will either pay for the difference or receive a refund.
  - If the price of the new instance class is higher than that of the original instance class, you need to pay for the price difference based on the used resource period.
  - If the price of the new instance class is lower than that of the original instance class, you will be refunded the difference based on the used resource period. The refund will be sent to your account. You can click **Billing Center** in the upper right corner of the console to view your account balance.
- For details, see **Product Pricing Details**.

Changing the Class of a Replica Set Instance

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the replica set instance and choose **More** > **Change Instance Class** in the **Operation** column.

Figure 7-23 Changing the class of a replica set instance

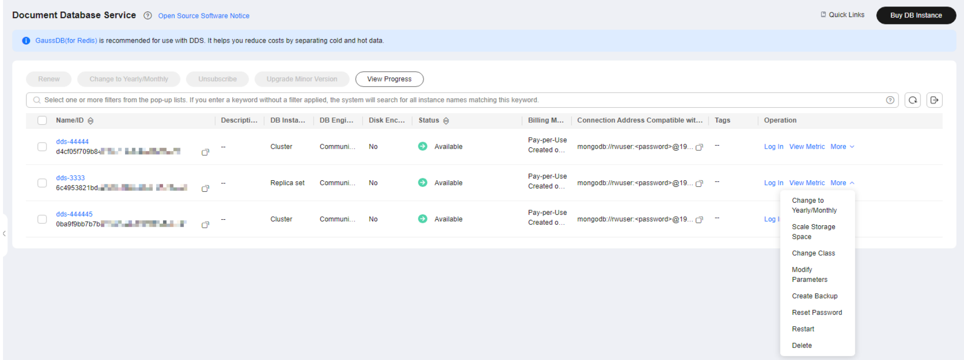

GaussDB(for Redis) is recommended for use with DDS. It helps you reduce costs by separating cold and hot data.

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all instance names matching this keyword.

Name/ID	Descripti...	DB Insta...	DB Engi...	Disk Enc...	Status	Billing M...	Connection Address Compatible wit...	Tags	Operation
dds-44444 d4cf05f709b84[redacted]	--	Cluster	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
dds-3333 6c4953821bd[redacted]	--	Replica set	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log In View Metric More
dds-444445 0ba9f9bb7b7b[redacted]	--	Cluster	Communi...	No	Available	Pay-per-Use Created o...	mongodb://rwuser:<password>@19...	--	Log I

Alternatively, on the **Instances** page, click the name of the replica set instance. In the **DB Information** area on the **Basic Information** page, click **Change Class** to the right of the **Node Class** field.

Figure 7-24 Changing the class of a replica set instance

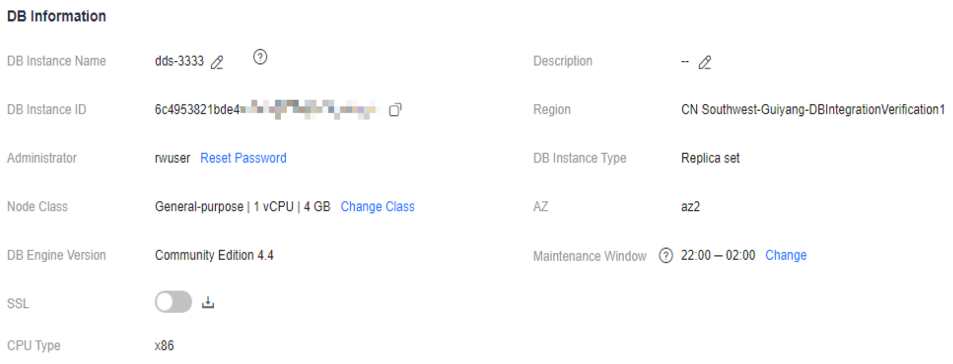

DB Information

DB Instance Name	dds-3333	Description	--
DB Instance ID	6c4953821bde4[redacted]	Region	CN Southwest-Guiyang-DBIntegrationVerification1
Administrator	rwuser Reset Password	DB Instance Type	Replica set
Node Class	General-purpose | 1 vCPU | 4 GB Change Class	AZ	az2
DB Engine Version	Community Edition 4.4	Maintenance Window	22:00 – 02:00 Change
SSL			
CPU Type	x86		

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

NOTE

The time required depends on the number of instance nodes whose class is to be changed. It takes about 5 to 10 minutes for each node. When the instance class is being changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted each time for up to 30s. Before the class change, learn about **Pre-check Items for Instance Class Change**. You are advised to change the class during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.

Figure 7-25 Changing the class of a replica set instance

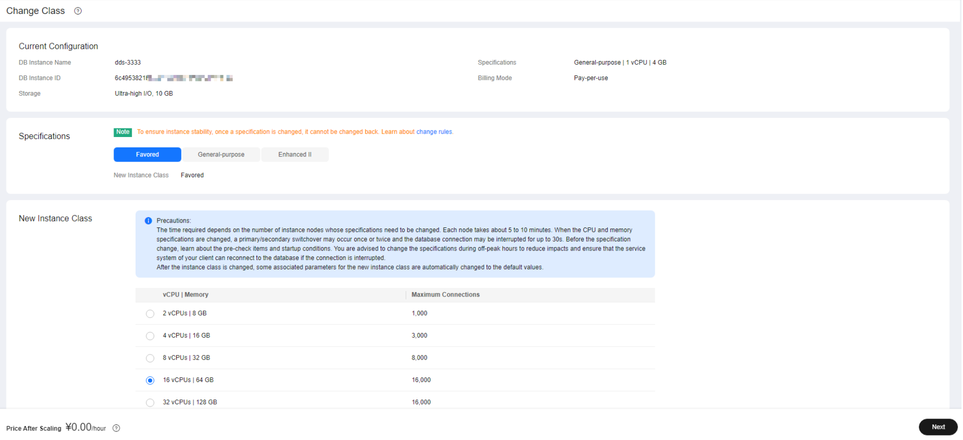

Change Class

Current Configuration

DB Instance Name	dds-3333	Specifications	General-purpose | 1 vCPU | 4 GB
DB Instance ID	6c4953821f[redacted]	Billing Mode	Pay-per-use
Storage	Ultra-high I/O, 10 GB		

Specifications

Note To ensure instance stability, once a specification is changed, it cannot be changed back. Learn about change rules.

New Instance Class Favored

New Instance Class

Precautions:
The time required depends on the number of instance nodes whose specifications need to be changed. Each node takes about 5 to 10 minutes. When the CPU and memory specifications are changed, a primary/secondary switchover may occur once or twice and the database connection may be interrupted for up to 30s. Before the specification change, learn about the pre-check items and startup conditions. You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
After the instance class is changed, some associated parameters for the new instance class are automatically changed to the default values.

	vCPU | Memory	Maximum Connections
○	2 vCPUs | 8 GB	1,000
○	4 vCPUs | 16 GB	3,000
○	8 vCPUs | 32 GB	8,000
●	16 vCPUs | 64 GB	16,000
○	32 vCPUs | 128 GB	16,000

Price After Scaling ¥0.00/hour

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

- If you need to modify your settings, click **Previous**.
- For pay-per-use instances

  If you do not need to modify your settings, click **Submit** to change the instance class. After the class is changed, you are still charged on an hourly basis.
- For yearly/monthly instances
  - If you intend to scale down the instance class, click **Submit**. The refund is automatically returned to your account.
  - If you intend to scale up the DB instance class, click **Pay Now**. The scaling starts only after the payment is successful.

Step 7 View the results.

- When the instance class is being changed, the status displayed in the **Status** column is **Changing instance class**. This process takes about 25 to 30 minutes.

  NOTE

  High database load increases the class change duration. You are advised to change the class during off-peak hours to reduce impacts.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.
- Go to the **Basic Information** page of the replica set instance you scaled up and check whether the scaling up is successful in the **DB Information** area.

----End

Changing the Class of a Read Replica

NOTE

Only whitelisted users can change the class of a read replica. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target replica set instance and click its name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **Read replicas** tab. Locate the read replica whose class you want to change, and click **Change Instance Class** in the **Operation** column.

Figure 7-26 Changing the class of a read replica

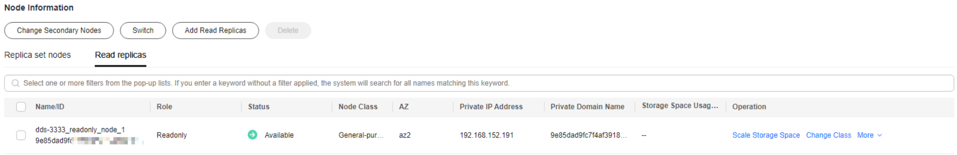

Node Information

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

□	Name/ID	Role	Status	Node Class	AZ	Private IP Address	Private Domain Name	Storage Space Usag...	Operation
□	dds-3333_readonly_node_1 9e85dad9f[redacted]	Readonly	Available	General-pur...	az2	192.168.152.191	9e85dad9fc7f4af3918...	--	Scale Storage Space Change Class More

Step 6 On the displayed page, select the required specifications and new class and click **Next**.

NOTE

When the class of a read replica is being changed, there is a possibility that database access requests using the read replica fail. Before the class change, learn about **Pre-check Items for Instance Class Change**. You are advised to change the class during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.

Figure 7-27 Changing the class of a read replica

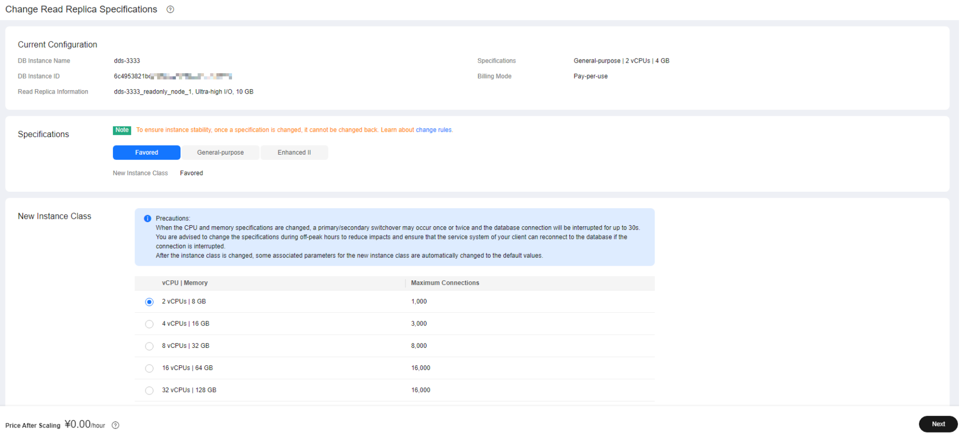

Change Read Replica Specifications

Current Configuration

DB Instance Name: dds-3333

DB Instance ID: 6c4953821b[redacted]

Read Replica Information: dds-3333_readonly_node_1, Ultra-high I/O, 10 GB

Specifications: General-purpose | 2 vCPUs | 4 GB

Billing Mode: Pay-per-use

Specifications

Note To ensure instance stability, once a specification is changed, it cannot be changed back. Learn about change rules.

New Instance Class: Favored

New Instance Class

Precautions:
When the CPU and memory specifications are changed, a primary/secondary switchover may occur once or twice and the database connection will be interrupted for up to 30s.
You are advised to change the specifications during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted.
After the instance class is changed, some associated parameters for the new instance class are automatically changed to the default values.

	vCPU | Memory	Maximum Connections
●	2 vCPUs | 8 GB	1,000
○	4 vCPUs | 16 GB	3,000
○	8 vCPUs | 32 GB	8,000
○	16 vCPUs | 64 GB	16,000
○	32 vCPUs | 128 GB	16,000

Price After Scaling ¥0.00/hour

Step 7 On the displayed page, confirm the class.

- If you need to modify your settings, click **Previous**.
- For pay-per-use instances

  If you do not need to modify your settings, click **Submit** to change the class. After the class is changed, you are still charged on an hourly basis.
- For yearly/monthly instances
  - If you intend to scale down the class, click **Submit**. The refund is automatically returned to your account.

  – If you intend to scale up the class, click **Pay Now**. The scaling starts only after the payment is successful.

Step 8 View the results.

- When the class is being changed, the status displayed in the **Status** column is **Changing instance class**. This process takes about 25 to 30 minutes.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.
- In the **Node Information** area on the **Basic Information** page, click the **Read replicas** tab. Locate the target read replica to view the new class.

----End

Reference

How Do I Solve the High CPU Usage Issue?

7.8.3 Changing a Single Node Instance Class

This section describes how to change the class of your single node instance.

Change Rules

Considering the stability and performance of DDS DB instances, you can change the DB instance class according to the rules listed in **Table 7-4**. Exercise caution when performing this operation.

Table 7-4 Change rules

Original Specification	Target Specification	Supported
General-purpose	General-purpose	√
	Enhanced	×
	Enhanced II	√
Enhanced	General-purpose	√
	Enhanced	×
	Enhanced II	√
Enhanced II	General-purpose	×
	Enhanced	×
	Enhanced II	√

NOTE

√ indicates that an item is supported, and × indicates that an item is not supported.

Precautions

- An instance cannot be deleted while its instance class is being changed.
- When the instance class is being changed, the database connection will be interrupted for 5 to 10 minutes. You are advised to change the class during off-peak hours to reduce impacts and ensure that the service system of your client can reconnect to the database if the connection is interrupted. After the restart is complete, the cached memory will be automatically cleared. The instance needs to be warmed up to prevent congestion during peak hours.
- After the class of a single node instance is changed, the system will change the value of **net.maxIncomingConnections** accordingly.
- Changing the class does not cause data loss.
- If you forcibly change the class of an abnormal node in a DB instance, services may be interrupted.

  NOTE

  To forcibly change the class of an abnormal node, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console.

Pre-check Items for Instance Class Change

- The DB instance is in the **Available** status.

Billing

- Instances billed on a pay-per-use basis are still billed based on the time used after the instance class is changed.
- If you change the class of a yearly/monthly instance, you will either pay for the difference or receive a refund.
  - If the price of the new instance class is higher than that of the original instance class, you need to pay for the price difference based on the used resource period.
  - If the price of the new instance class is lower than that of the original instance class, you will be refunded the difference based on the used resource period. The refund will be sent to your account. You can click **Billing Center** in the upper right corner of the console to view your account balance.
- For details, see **Product Pricing Details**.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the single node instance and choose **More** > **Change Instance Class** in the **Operation** column.

Alternatively, on the **Instances** page, click the name of the single node instance. In the **DB Information** area on the **Basic Information** page, click **Change** to the right of the **Node Class** field.

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

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

- If you need to modify your settings, click **Previous**.
- For pay-per-use instances

  If you do not need to modify your settings, click **Submit** to change the instance class. After the specifications are changed, you are still charged on an hourly basis.
- For yearly/monthly instances
  - If you intend to scale down the instance class, click **Submit**. The refund is automatically returned to your account.
  - If you intend to scale up the instance class, click **Pay Now**. The scaling starts only after the payment is successful.

Step 7 View the results.

- When the instance class is being changed, the status displayed in the **Status** column is **Changing instance class**. This process takes about 10 minutes.

  > NOTE
  >
  > High database load increases the specification change duration. You are advised to change the specifications during off-peak hours to reduce impacts.
- In the upper right corner of the instance list, click ↻ to refresh the list. The instance status changes to **Available**.
- Go to the **Basic Information** page of the single node you scaled up and check whether the scaling process is successful in the **Configuration** area.

----End

Reference

How Do I Solve the High CPU Usage Issue?

7.9 Changing Cluster Instance Nodes

7.9.1 Adding Cluster Instance Nodes

As service data increases, the number of current database nodes cannot meet the service requirements. In this case, you can add more nodes to the instance.

Precautions

- To add nodes, instance status must be **Available**, **Deleting backup**, or **Checking restoration**.
- Nodes cannot be added to a DB instance that is being backed up.

- A DB instance cannot be deleted while nodes are being added.
- An instance node can be added within 5 minutes. The time required depends on the number of nodes to be added.
- Adding nodes does not affect cluster services.
- When adding a shard node for a cluster DB instance, ensure that the node class is greater than or equal to the highest class of a shard in the instance.

Pricing Details

- A pay-per-use instance is still billed on an hourly basis after new nodes are added.
- If you add nodes to a yearly/monthly instance, you will pay price difference or get a refund.
- For details, see **Product Pricing Details**.

Adding dds mongos Nodes

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 On the **dds mongos** tab in the **Node Information** area, click **Add dds mongos**.

Figure 7-28 Node information

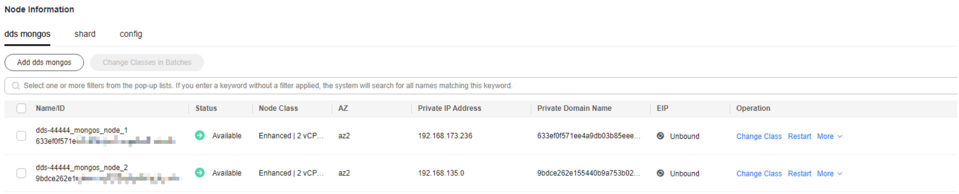

Node Information

dds mongos | shard | config

Add dds mongos | Change Classes in Batches

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

□	Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name	EIP	Operation
□	dds-44444_mongos_node_1 633ef0f571e[redacted]	Available	Enhanced | 2 vCP...	az2	192.168.173.236	633ef0f571ee4a9db03b85eee...	Unbound	Change Class Restart More
□	dds-44444_mongos_node_2 9bdce262e1[redacted]	Available	Enhanced | 2 vCP...	az2	192.168.135.0	9bdce262e155440b9a753b02...	Unbound	Change Class Restart More

Step 6 On the displayed page, specify **Node Class**, **Nodes**, and **Parameter Template** and click **Next**.

Figure 7-29 Adding dds mongos nodes

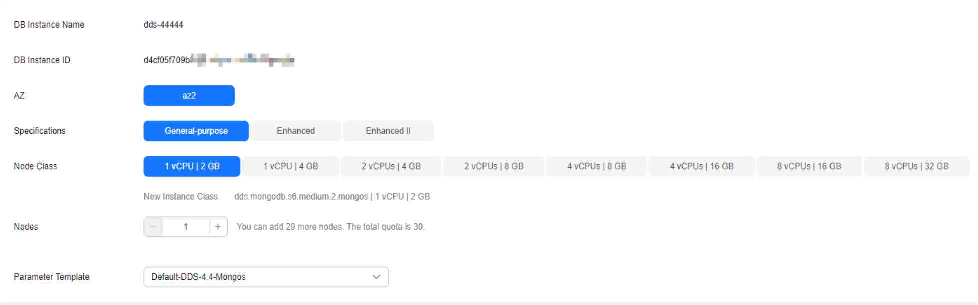

DB Instance Name	dds-44444
DB Instance ID	d4cf05f709b[redacted]
AZ	az2
Specifications	General-purpose / Enhanced / Enhanced II
Node Class	1 vCPU | 2 GB / 1 vCPU | 4 GB / 2 vCPUs | 4 GB / 2 vCPUs | 8 GB / 4 vCPUs | 8 GB / 4 vCPUs | 16 GB / 8 vCPUs | 16 GB / 8 vCPUs | 32 GB
	New Instance Class dds.mongodb.s6.medium.2.mongos | 1 vCPU | 2 GB
Nodes	− 1 + You can add 29 more nodes. The total quota is 30.
Parameter Template	Default-DDS-4.4-Mongos

A Community Edition cluster instance supports up to 32 dds mongos nodes.

Step 7 On the displayed page, confirm the node configuration information.

- Yearly/Monthly
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify your settings, click **Submit** to go to the payment page and complete the payment.
- Pay-per-use
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify your settings, click **Submit** to add the nodes.

Step 8 View the results.

- This process takes about 10 to 15 minutes. During that time, the status of the DB instance in the instance list is **Adding node**.
- In the upper right corner of the DB instance list, click ⟳ to refresh the list. The instance status changes to **Available**.
- On the **dds mongos** tab in the **Node Information** area, view the information about the node you added.
- If the dds mongos nodes fail to be added, you can revert them in batches or delete them one by one. For details, see section **Reverting Cluster Instance Nodes**.

----End

Adding shard Nodes

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instance Management** page, click the target cluster instance.

Step 5 On the **shard** tab in the **Node Information** area, click **Add shard**.

Figure 7-30 Node information

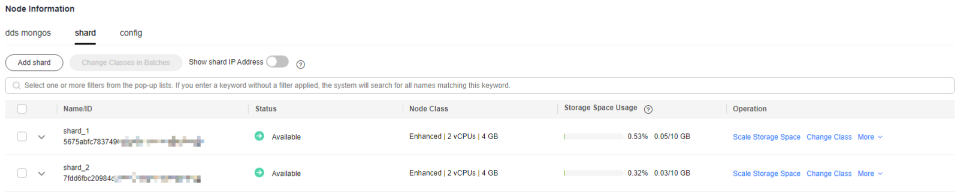

Node Information

dds mongos | shard | config

Add shard | Change Classes in Batches | Show shard IP Address

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

□	Name/ID	Status	Node Class	Storage Space Usage	Operation
□	shard_1 5675abfc783749[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	0.53% 0.05/10 GB	Scale Storage Space Change Class More
□	shard_2 7fdd6fbc20984c[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	0.32% 0.03/10 GB	Scale Storage Space Change Class More

Step 6 Specify **Node Class**, **Storage Space**, **Nodes**, and **Parameter Template** and click **Next**.

Figure 7-31 Adding shard nodes

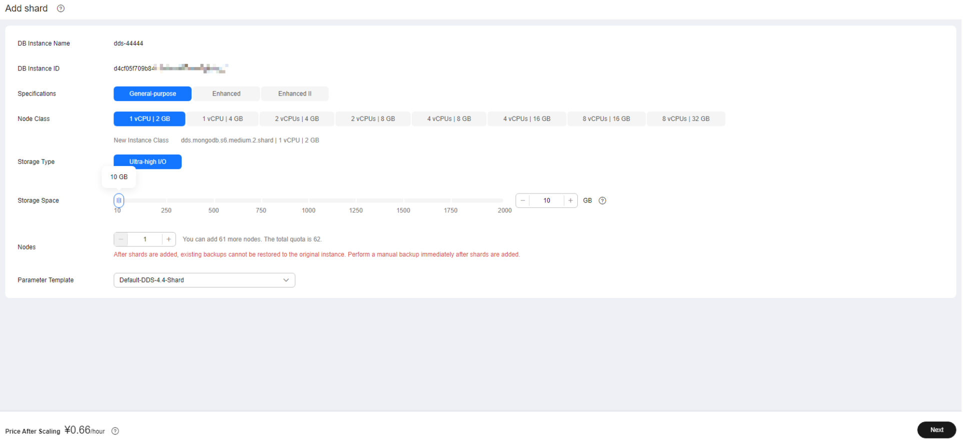

Add shard

DB Instance Name: dds-44444

DB Instance ID: d4cf05f709b84[redacted]

Specifications: General-purpose

Node Class: 1 vCPU | 2 GB

New Instance Class: dds.mongodb.s6.medium.2.shard | 1 vCPU | 2 GB

Storage Type: Ultra-high I/O

Storage Space: 10 GB

Nodes: 1 — You can add 61 more nodes. The total quota is 62.

After shards are added, existing backups cannot be restored to the original instance. Perform a manual backup immediately after shards are added.

Parameter Template: Default-DDS-4.4-Shard

Price After Scaling ¥0.66/hour

- The storage space you applied for will include the system overhead required for inode, reserved block, and database operation. The storage space must be a multiple of 10.
- A Community Edition cluster instance supports up to 32 shard nodes.

Step 7 On the displayed page, confirm the node configuration information.

- Yearly/Monthly
  – If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  – If you do not need to modify your settings, click **Submit** to go to the payment page and complete the payment.
- Pay-per-use
  – If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  – If you do not need to modify your settings, click **Submit** to add the nodes.

Step 8 View the results.

- This process takes about 10 to 15 minutes. During that time, the status of the DB instance in the instance list is **Adding node**.
- In the upper right corner of the DB instance list, click ⟳ to refresh the list. The instance status changes to **Available**.
- On the **shard** tab in the **Node Information** area, view the information about the node you added.
- If shard addition fails, you can roll back the operation in batches or delete shards one by one. For details, see **Reverting Cluster Instance Nodes**.

----End

7.9.2 Adding Read Replicas to a Cluster Instance

Read replicas are used to enhance read capabilities and reduce the read pressure on primary nodes. After a DDS cluster instance is created, you can create read replicas based on service requirements.

Constraints

- You can add nodes only when your account balance is greater than or equal to $0 USD. To use this function, contact customer service to apply for the required permission.
- The cluster instance version must be 3.4.
- Nodes cannot be added to an instance that is being backed up.
- An instance cannot be deleted when one or more nodes are being added.
- The synchronization delay cannot be set. The default value is **0**.

Precautions

- A maximum of five read replicas can be added to a shard node.
- You can add read replicas to only one shard at a time.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **shard** tab, locate a target shard node, and choose **More** > **Add Read Replicas** in the **Operation** column.

Figure 7-32 Node information

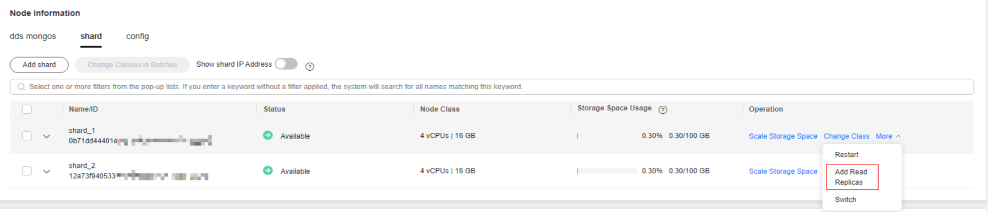
Node Information
dds mongos
shard
config
Add shard
Change Classes in Batches
Show shard IP Address
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.
Restart
Add Read Replicas
Switch

Name/ID	Status	Node Class	Storage Space Usage	Operation
shard_1 0b71dd44401e[redacted]	Available	4 vCPUs | 16 GB	0.30% 0.30/100 GB	Scale Storage Space Change Class More
shard_2 12a73f940533[redacted]	Available	4 vCPUs | 16 GB	0.30% 0.30/100 GB	Scale Storage Space

Step 6 On the displayed page, specify **Node Class**, **Nodes**, and **Parameter Template** and click **Next**.

Figure 7-33 Adding read replicas

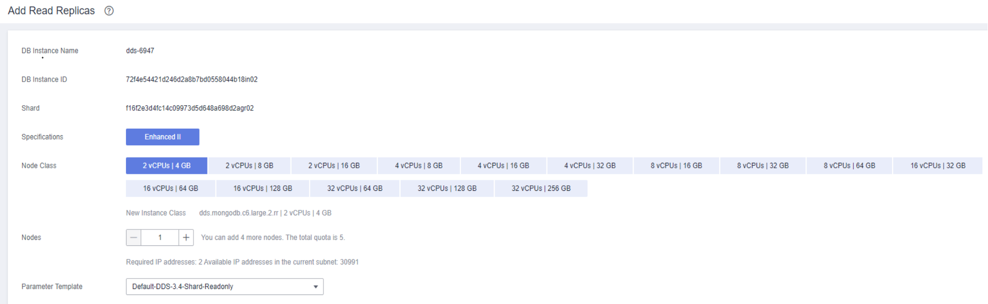

DB Instance Name: dds-6947

DB Instance ID: 72f4e54421d246d2a8b7bd0558044b18in02

Shard: f16f2e3d4fc14c09973d5d648a698d2agr02

Specifications: Enhanced II

Node Class: 2 vCPUs | 4 GB

New Instance Class dds.mongodb.c6.large.2.rr | 2 vCPUs | 4 GB

Nodes: 1 You can add 4 more nodes. The total quota is 5.

Required IP addresses: 2 Available IP addresses in the current subnet: 30991

Parameter Template: Default-DDS-3.4-Shard-Readonly

Table 7-5 Parameter description

Parameter	Description
Read Replica Parameter Template	The parameters that apply to the read replicas of a cluster instance. After a node are created, you can change the parameter template of the node to bring out the best performance. **NOTE** Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Step 7 On the displayed page, confirm the node configuration information.

- If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
- If you do not need to modify your settings, click **Submit** to add nodes.

Step 8 View the results.

- When nodes are added, the status of the instance is **Adding read replicas**. The entire process takes about 15 minutes.
- In the **Node Information** area, view the information about the nodes you added.
- Choose **More** > **View Delay** in the **Operation** column to view the delay of the current node.

----End

7.9.3 Manually Switching the Primary and Secondary Nodes of a Cluster

A cluster instance consists of a config node, and multiple dds mongos and shard nodes. Each shard or config is deployed as a three-node replica set including primary, secondary, and hidden nodes. Primary and secondary nodes do not provide IP addresses for external access. Hidden nodes are only used for backing up data. When a primary node becomes faulty, the system automatically selects a new primary node to ensure high availability. DDS supports primary/secondary switchovers for scenarios such as disaster recovery.

Precautions

- To perform a switchover, the instance status needs to be **Available**, **Abnormal**, **Changing to yearly/monthly**, **Changing a security group**, or **Heavy load**.
- The database connection may be interrupted during the switchover. Ensure that your client supports reconnection.
- The longer the delay for primary/secondary synchronization, the more time is needed for a primary/secondary switchover. If the primary to secondary synchronization delay exceeds 300s, primary/secondary switchover is not supported. For details about the synchronization delay, see **What Is the Time Delay for Primary/Secondary Synchronization in a Replica Set?**

Performing a Primary/Secondary Switchover for a Config Node

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **config** tab and click **Switch** in the **Operation** column.

Figure 7-34 Primary/Secondary switchover

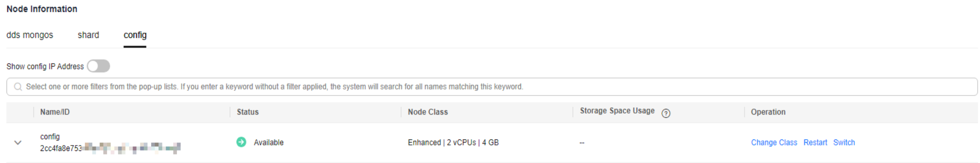

Node Information

dds mongos shard config

Show config IP Address

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Status	Node Class	Storage Space Usage	Operation
config 2cc4fa8e753[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	--	Change Class Restart Switch

Step 6 In the displayed dialog box, click **Yes**.

Figure 7-35 Performing a primary/secondary switchover for a config node

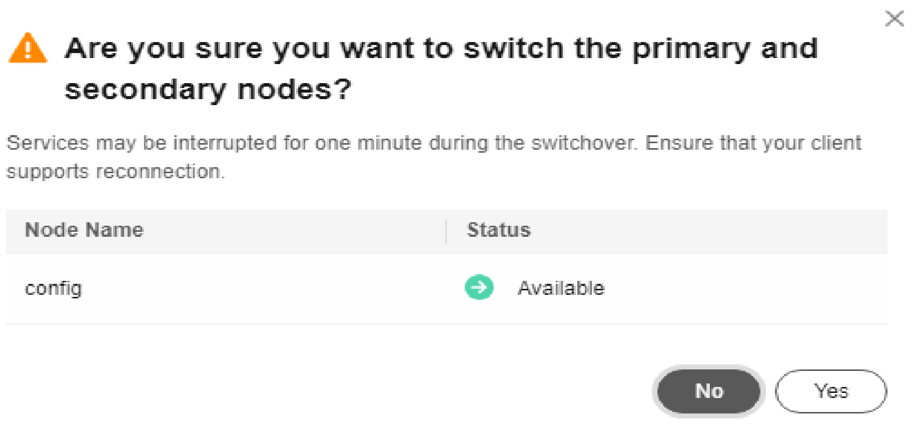

Are you sure you want to switch the primary and secondary nodes?

Services may be interrupted for one minute during the switchover. Ensure that your client supports reconnection.

Node Name	Status
config	Available

Step 7 Check the result.

- During the switchover process, the DB instance status changes to **Switchover in progress**. After the switchover is complete, the status is restored to **Available**.
- In the **Node Information** area, you can view the switchover result.
- After the switchover, the previous primary node becomes the secondary node. You need to reconnect to the primary node. For details, see **Connecting to a DB Instance**.

----End

Performing a Primary/Secondary Switchover for a Shard Node

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **shard** tab, locate a target shard node, and choose **More** > **Switch** in the **Operation** column.

Figure 7-36 Primary/Secondary switchover

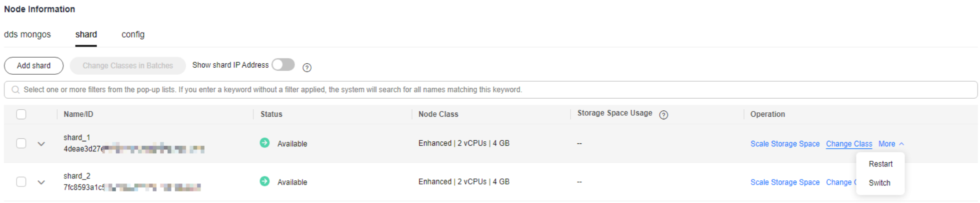

Node Information

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Status	Node Class	Storage Space Usage	Operation
shard_1 4deae3d27[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	--	Scale Storage Space Change Class More
shard_2 7fc8593a1c[redacted]	Available	Enhanced | 2 vCPUs | 4 GB	--	Scale Storage Space Change C

Step 6 In the displayed dialog box, click **Yes**.

Figure 7-37 Performing a primary/secondary switchover for a shard node

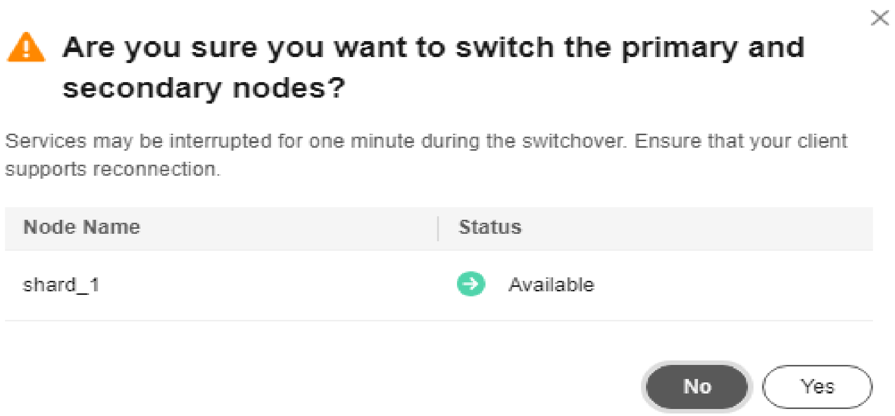

Are you sure you want to switch the primary and secondary nodes?

Services may be interrupted for one minute during the switchover. Ensure that your client supports reconnection.

Node Name	Status
shard_1	Available

Step 7 Check the result.

- During the switchover process, the DB instance status changes to **Switchover in progress**. After the switchover is complete, the status is restored to **Available**.
- In the **Node Information** area, you can view the switchover result.
- After the switchover, the previous primary node becomes the secondary node. You need to reconnect to the primary node. For details, see **Connecting to a DB Instance**.

----End

Forcibly Promoting a Secondary Node to the Primary

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click the **config** or **shard** tab, locate the node whose role is **Secondary**, and click **Promote to Primary**.

Figure 7-38 Promote to Primary

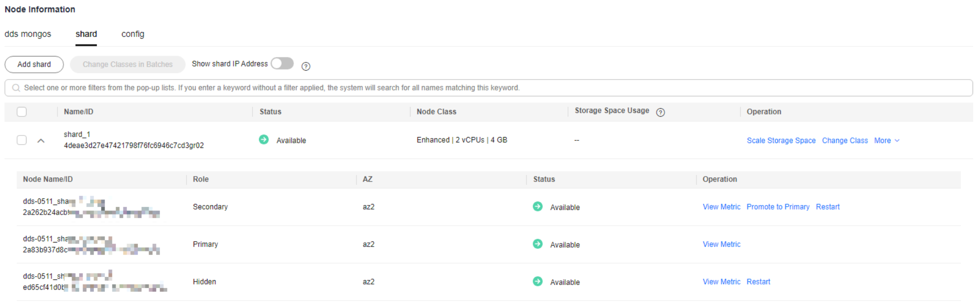

Node Information

dds mongos | shard | config

Add shard | Change Classes in Batches | Show shard IP Address

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Status	Node Class	Storage Space Usage	Operation
shard_1 4deae3d27e47421798f76fc6946c7cd3gr02	Available	Enhanced | 2 vCPUs | 4 GB	--	Scale Storage Space Change Class More

Node Name/ID	Role	AZ	Status	Operation
dds-0511_sha[redacted] 2a262b24acb[redacted]	Secondary	az2	Available	View Metric Promote to Primary Restart
dds-0511_sha[redacted] 2a83b937d8c[redacted]	Primary	az2	Available	View Metric
dds-0511_sh[redacted] ed65cf41d0[redacted]	Hidden	az2	Available	View Metric Restart

Step 6 In the displayed dialog box, click **Yes**.

Figure 7-39 Forcibly promoting a secondary node to the primary

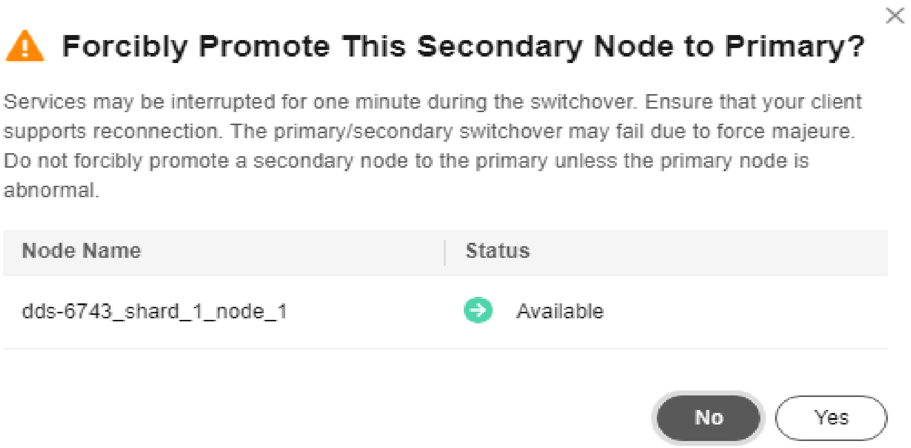

Forcibly Promote This Secondary Node to Primary?

Services may be interrupted for one minute during the switchover. Ensure that your client supports reconnection. The primary/secondary switchover may fail due to force majeure. Do not forcibly promote a secondary node to the primary unless the primary node is abnormal.

Node Name	Status
dds-6743_shard_1_node_1	Available

Step 7 Check the result.

- In the **Node Information** area on the **Basic Information** page, you can view the result.

----End

7.9.4 Reverting Cluster Instance Nodes

This section describes how to roll back a failed node addition.

Reverting Nodes in Batches

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the cluster instance to which nodes fail to be added and choose **More** > **Revert** in the **Operation** column.

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

During the rollback, the instance status is **Deleting node**. This process takes about 1 to 3 minutes.

----End

Deleting a Single Node

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance to which the node fails to be added.

Step 5 In the **Node Information** area on the **Basic Information** tab, click the **dds mongos** or **shard** tab, locate the dds mongos node, shard node, or read replica that fails to be added, and click **Delete**.

Step 6 In the displayed dialog box, click **Yes**.

During deletion, the node status is **Deleting node**. This process takes about 1 to 3 minutes.

----End

7.10 Changing Replica Set Instance Nodes

7.10.1 Adding Replica Set Instance Nodes

DDS allows you to scale out a three-node replica set instance to up to five or even seven nodes. All newly added nodes are secondary nodes and support primary/secondary switchovers, improving data reliability.

NOTE

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Precautions

- To add nodes, instance status must be **Available**, **Deleting backup**, or **Checking restoration**.
- A DB instance cannot be deleted while nodes are being added.
- If there are any newly added standby nodes, they will be unable to participate in this switchover. When you add a new standby node, the HA connection address needs to be reconfigured, and the new node is frozen for 12 hours.
- When instance nodes are being added, the DB instance may be intermittently disconnected once or twice for up to 30s each time.
- Nodes cannot be manually deleted.

Pricing Details

- A pay-per-use instance is still billed on an hourly basis after new nodes are added.
- If you add nodes to a yearly/monthly instance, you will pay price difference or get a refund.
- For details, see **Product Pricing Details**.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, click **Change Secondary Nodes**.

Figure 7-40 Basic information

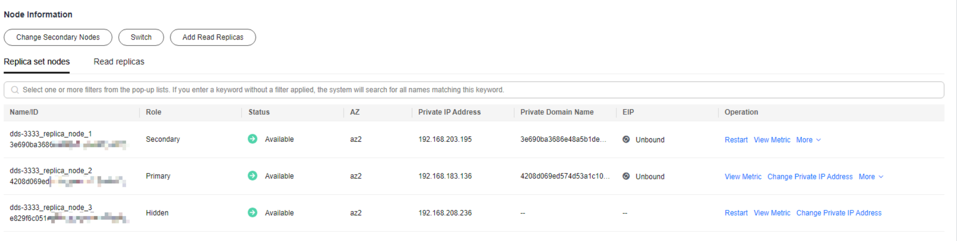
Node Information
Change Secondary Nodes
Switch
Add Read Replicas
Replica set nodes
Read replicas
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Role	Status	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-3333_replica_node_1 3e690ba3686[redacted]	Secondary	Available	az2	192.168.203.195	3e690ba3686e48a5b1de...	Unbound	Restart View Metric More
dds-3333_replica_node_2 4208d069ed[redacted]	Primary	Available	az2	192.168.183.136	4208d069ed574d53a1c10...	Unbound	View Metric Change Private IP Address More
dds-3333_replica_node_3 e829f6c051[redacted]	Hidden	Available	az2	192.168.208.236	--	--	Restart View Metric Change Private IP Address

Step 6 Specify **Total Nodes** and click **Next**.

Figure 7-41 Selecting the number of nodes

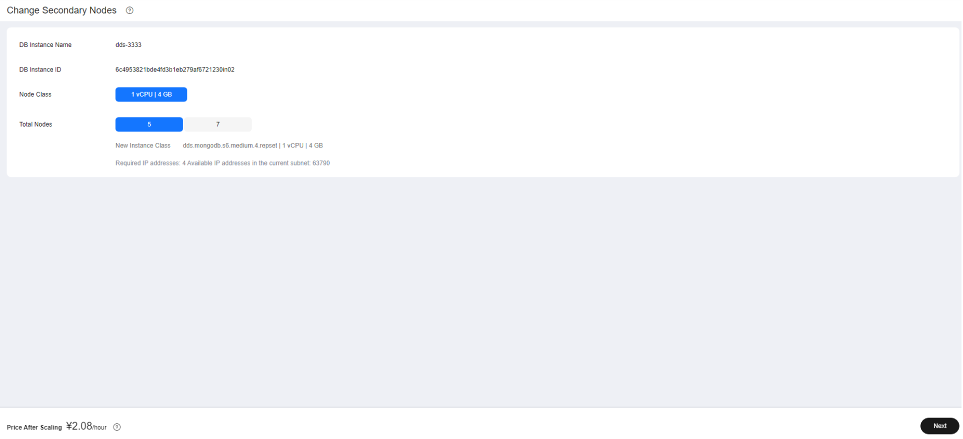

Change Secondary Nodes

DB Instance Name dds-3333

DB Instance ID 6c4953821bde4fd3b1eb279af6721230in02

Node Class 1 vCPU | 4 GB

Total Nodes 5

New Instance Class dds.mongodb.s6.medium.4.repset | 1 vCPU | 4 GB

Required IP addresses: 4 Available IP addresses in the current subnet: 63790

Price After Scaling ¥2.08/hour

You can add five or seven nodes.

Step 7 On the displayed page, confirm the node configuration information.

- For yearly/monthly DB instances
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify your settings, click **Submit** to go to the payment page and complete the payment.
- For pay-per-use DB instances
  - If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
  - If you do not need to modify your settings, click **Submit** to add the nodes.

Step 8 View the result of adding nodes.

- When a node is being added, the status of the instance is **Adding node**. The entire process takes about 15 minutes.
- In the **Node Information** area, view the information about the nodes you added.

----End

7.10.2 Adding Read Replicas to a Replica Set Instance

Read replicas enhance read capabilities and reduce load on your instances. After a DDS replica set instance is created, you can create read replicas based on service requirements. To connect to a read replica, see **Connecting to a Read Replica Using Mongo Shell**.

Constraints

- To use this function, contact customer service to apply for the required permission.
- The version of a replica set instance must be 3.4, 4.0, 4.2, 4.4 or 5.0.
- Nodes cannot be added to an instance that is being backed up.

- An instance cannot be deleted when one or more nodes are being added.
- When read replicas are being added, the DB instance may be intermittently disconnected once or twice for up to 30s each time.

Precautions

- A maximum of five read replicas can be added to a replica set instance.

Procedure

Step 1 Log in to the management console.

Step 2 Click ⚲ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance.

Step 5 In the **Node Information** area on the **Basic Information** page, click **Add Read Replicas**.

Figure 7-42 Creating read replicas

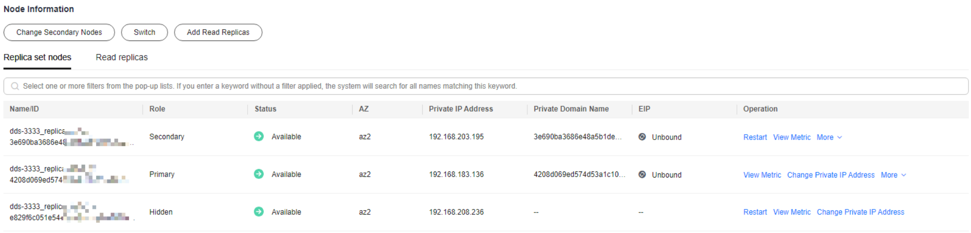

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Role	Status	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-3333_replica[redacted] 3e690ba3686e4[redacted]	Secondary	Available	az2	192.168.203.195	3e690ba3686e48a5b1de...	Unbound	Restart View Metric More
dds-3333_replic[redacted] 4208d069ed574[redacted]	Primary	Available	az2	192.168.183.136	4208d069ed574d53a1c10...	Unbound	View Metric Change Private IP Address More
dds-3333_replic[redacted] e829f6c051e54[redacted]	Hidden	Available	az2	192.168.208.236	--	--	Restart View Metric Change Private IP Address

Step 6 On the **Add Read Replicas** page, specify **Specifications**, **Node Class**, **Nodes**, **Parameter Template**, and **Delay**, and click **Next**.

Figure 7-43 Creating read replicas

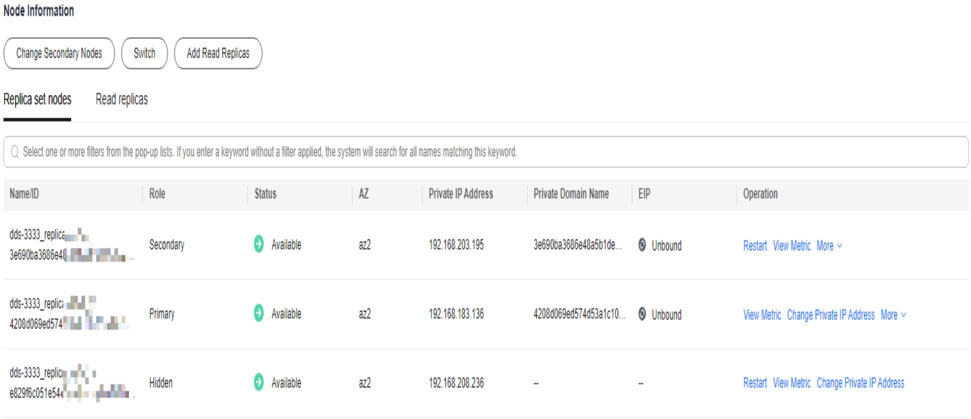

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Role	Status	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-3333_replica[redacted] 3e690ba3686e4[redacted]	Secondary	Available	az2	192.168.203.195	3e690ba3686e48a5b1de...	Unbound	Restart View Metric More
dds-3333_replic[redacted] 4208d069ed574[redacted]	Primary	Available	az2	192.168.183.136	4208d069ed574d53a1c10...	Unbound	View Metric Change Private IP Address More
dds-3333_replic[redacted] e829f6c051e54[redacted]	Hidden	Available	az2	192.168.208.236	--	--	Restart View Metric Change Private IP Address

Table 7-6 Parameter description

Parameter	Description
Read Replica Parameter Template	The parameters that apply to the read replicas of a replica set instance. After a node are created, you can change the parameter template of the node to bring out the best performance. **NOTE** Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Step 7 On the displayed page, confirm the node configuration information.

- If you need to modify your settings, click **Previous** to go back to the page for you to specify details.
- If you do not need to modify your settings, click **Submit** to add nodes.

Step 8 View the results.

- When nodes are added, the status of the instance is **Adding read replicas**. The entire process takes about 15 minutes.
- In the **Node Information** area, view the information about the nodes you added.
- Choose **More** > **View Delay** in the **Operation** column to view the delay of the current node.

----End

7.10.3 Manually Switching the Primary and Secondary Nodes of a Replica Set

A replica set consists of the primary node, secondary node, and hidden node. Primary and secondary nodes allow access from external services by providing IP addresses. Hidden nodes are only used for backing up data. When a primary node becomes faulty, the system automatically selects a new primary node to ensure high availability. DDS supports primary/secondary switchovers for scenarios such as disaster recovery.

Precautions

- To perform a switchover, the instance status needs to be **Available**, **Changing to yearly/monthly**, and **Changing a security group**.
- The database connection may be interrupted during the switchover. Ensure that your client supports reconnection.
- If there are any newly added secondary nodes, they will be unable to participate in this switchover. When you add a new secondary node, the HA connection address needs to be reconfigured, and the new node is frozen for 12 hours.
- A primary/secondary switchover can be performed only when the DB instance is available.

- The longer the delay for primary/secondary synchronization, the more time is needed for a primary/secondary switchover. If the primary to secondary synchronization delay exceeds 300s, primary/secondary switchover is not supported. For details about the synchronization delay, see **What Is the Time Delay for Primary/Secondary Synchronization in a Replica Set?**

Performing a Primary/Secondary Switchover

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance.

Step 5 In the **Node Information** area on the **Basic Information** page, click **Switch**.

Figure 7-44 Primary/Secondary switchover

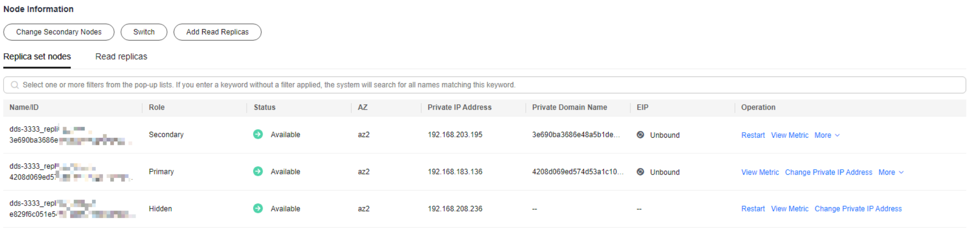

Node Information
Change Secondary Nodes
Switch
Add Read Replicas
Replica set nodes
Read replicas
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Role	Status	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-3333_repli [redacted] 3e690ba3686e [redacted]	Secondary	Available	az2	192.168.203.195	3e690ba3686e48a5b1de...	Unbound	Restart View Metric More
dds-3333_repl [redacted] 4208d069ed5 [redacted]	Primary	Available	az2	192.168.183.136	4208d069ed574d53a1c10...	Unbound	View Metric Change Private IP Address More
dds-3333_repl [redacted] e829f6c051e5 [redacted]	Hidden	Available	az2	192.168.208.236	--	--	Restart View Metric Change Private IP Address

Step 6 In the displayed dialog box, click **Yes**.

Step 7 Check the result.

- During the switchover process, the DB instance status changes to **Switchover in progress**. After the switchover is complete, the status is restored to **Available**.
- In the **Node Information** area, you can view the switchover result.
- After the switchover, the previous primary node becomes the secondary node. You need to reconnect to the primary node. For details, see **Connecting to a DB Instance**.

----End

Forcibly Promoting a Secondary Node to the Primary

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target replica set instance and click its name.

Step 5 In the **Node Information** area on the **Basic Information** page, locate a target node whose role is **Secondary** and click **Promote to Primary** in the **Operation** column.

Figure 7-45 Promote to Primary

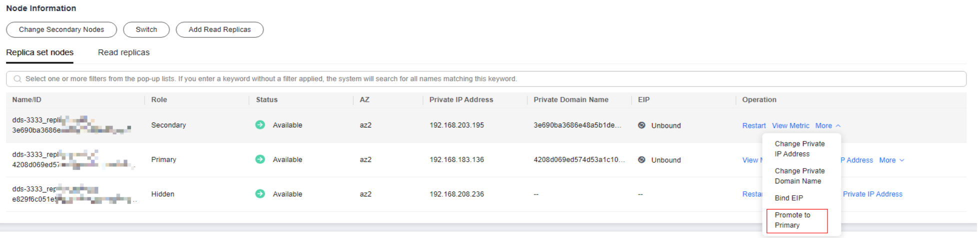
Node Information
Change Secondary Nodes
Switch
Add Read Replicas
Replica set nodes
Read replicas
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.
Change Private IP Address
Change Private Domain Name
Bind EIP
Promote to Primary

Name/ID	Role	Status	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-3333_repl[redacted] 3e690ba3686e[redacted]	Secondary	Available	az2	192.168.203.195	3e690ba3686e48a5b1de...	Unbound	Restart View Metric More
dds-3333_repl[redacted] 4208d069ed57[redacted]	Primary	Available	az2	192.168.183.136	4208d069ed574d53a1c10...	Unbound	View P Address More
dds-3333_rep[redacted] e829f6c051e[redacted]	Hidden	Available	az2	192.168.208.236	--	--	Resta Private IP Address

Step 6 In the displayed dialog box, click **Yes**.

Figure 7-46 Forcibly promoting a secondary node to the primary

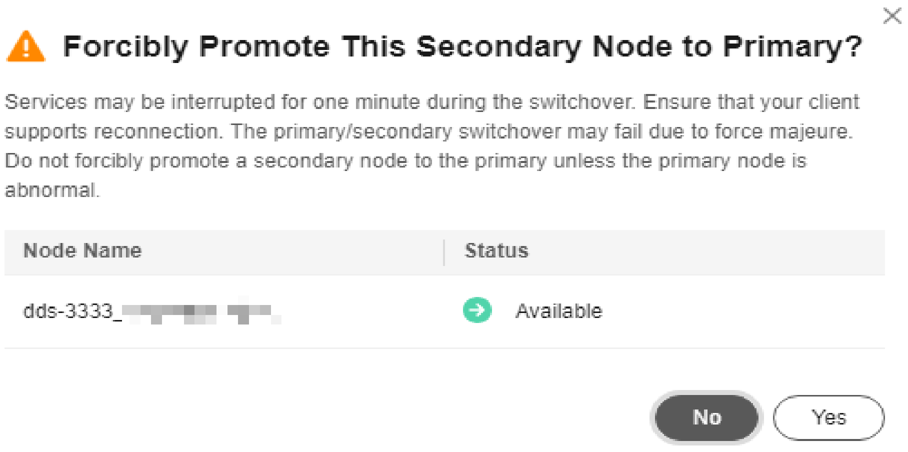

Forcibly Promote This Secondary Node to Primary?

Services may be interrupted for one minute during the switchover. Ensure that your client supports reconnection. The primary/secondary switchover may fail due to force majeure. Do not forcibly promote a secondary node to the primary unless the primary node is abnormal.

Node Name	Status
dds-3333_[redacted]	Available

Step 7 Check the result.

- In the **Node Information** area on the **Basic Information** page, you can view the result.

----End

7.10.4 Deleting Replica Set Instance Nodes

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

NOTE

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Precautions

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

- If you enable operation protection, two-factor authentication is required for sensitive operations to secure your account and cloud products. For details about how to enable operation protection, see *Identity and Access Management User Guide*.
- Nodes cannot be deleted from instances that have abnormal nodes.

Procedure

Step 1 Log in to the console.

Step 2 In the service list, choose **Databases** > **Document Database Service**.

Step 3 On the **Instances** page, locate a target DB instance and click its name.

Step 4 In the **Node Information** area on the **Basic Information** page, click **Change Secondary Nodes**.

Figure 7-47 Node information

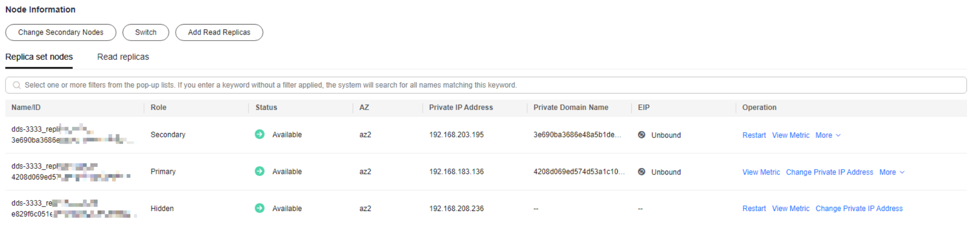
Node Information
Change Secondary Nodes
Switch
Add Read Replicas
Replica set nodes
Read replicas
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

Name/ID	Role	Status	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-3333_repl[redacted] 3e690ba3686e[redacted]	Secondary	Available	az2	192.168.203.195	3e690ba3686e48a5b1de...	Unbound	Restart View Metric More
dds-3333_repl[redacted] 4208d069ed5[redacted]	Primary	Available	az2	192.168.183.136	4208d069ed574d53a1c10...	Unbound	View Metric Change Private IP Address More
dds-3333_re[redacted] e829f6c051e[redacted]	Hidden	Available	az2	192.168.208.236	--	--	Restart View Metric Change Private IP Address

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

Figure 7-48 Selecting the number of nodes

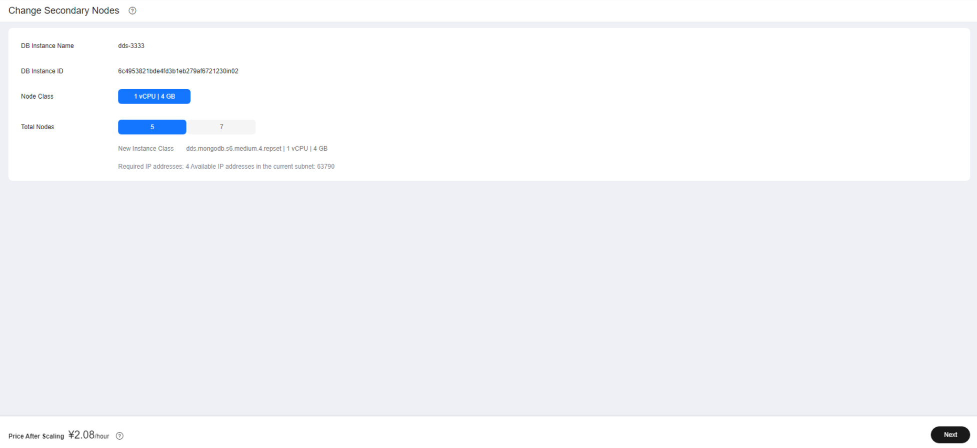

DB Instance Name: dds-3333

DB Instance ID: 6c4953821bde4fd3b1eb279af6721230in02

Node Class: 1 vCPU | 4 GB

Total Nodes: 5 7

New Instance Class dds.mongodb.s6.medium.4.repset | 1 vCPU | 4 GB

Required IP addresses: 4 Available IP addresses in the current subnet: 63790

Price After Scaling ¥2.08/hour

NOTE

When you delete nodes, the number of selected nodes should be less than the number of current nodes. For example, if the number of current instance nodes is 5, select **3** when deleting nodes.

Step 6 Click **Submit**.

Step 7 If you have enabled operation protection, click **Start Verification** in the **Delete Node** dialog box. On the displayed page, click **Send Code**, enter the verification code, and click **Verify**. The page is closed automatically.

----End

7.10.5 Deleting Read Replicas from a Replica Set Instance

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

NOTE

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Precautions

- Deleted read replicas cannot be restored. Exercise caution when performing this operation.
- If you have enabled operation protection, two-factor authentication is required for sensitive operations to secure your account and cloud products. For details about how to enable operation protection, see the *Identity and Access Management User Guide*.
- Read replicas cannot be deleted from instances that have abnormal nodes.

Procedure

Step 1 Log in to the management console.

Step 2 In the service list, choose **Databases** > **Document Database Service**.

Step 3 On the **Instances** page, locate a target DB instance and click its name.

Step 4 In the **Node Information** area on the **Basic Information** page, click the **Read replicas** tab.

- For yearly/monthly instances:
  - Locate a target read replica and choose **More** > **Delete** in the **Operation** column.

Figure 7-49 Selecting a read replica in the yearly/monthly instance

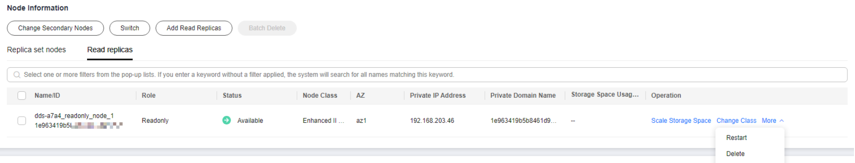

Node Information
Change Secondary Nodes
Switch
Add Read Replicas
Batch Delete
Replica set nodes
Read replicas
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.
Restart
Delete

Name/ID	Role	Status	Node Class	AZ	Private IP Address	Private Domain Name	Storage Space Usag...	Operation
dds-a7a4_readonly_node_1 1e963419b5l[redacted]	Readonly	Available	Enhanced II ...	az1	192.168.203.46	1e963419b5b8461d9...	--	Scale Storage Space Change Class More

  - Click **Yes**.

Figure 7-50 Deleting the read replica from the yearly/monthly instance

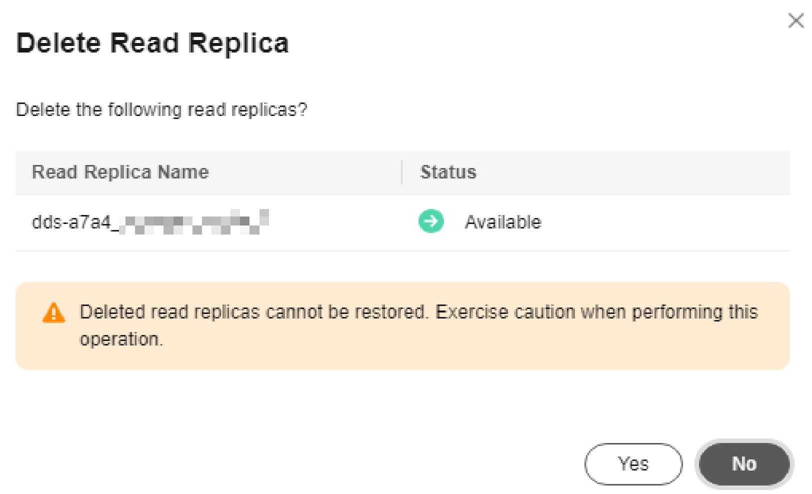

Delete Read Replica

Delete the following read replicas?

Read Replica Name	Status
dds-a7a4_[redacted]	Available

Deleted read replicas cannot be restored. Exercise caution when performing this operation.

- Select the check box before **I understand that a handling fee will be charged for this unsubscription** and click **OK**.

Figure 7-51 Confirming the deletion

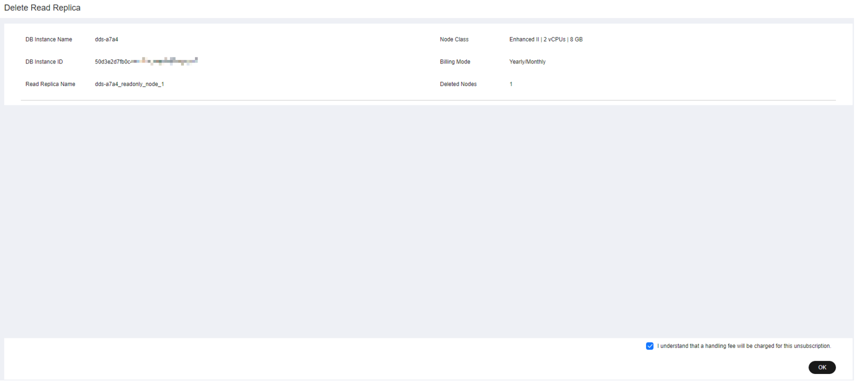

Delete Read Replica

DB Instance Name	dds-a7a4	Node Class	Enhanced II | 2 vCPUs | 8 GB
DB Instance ID	50d3e2d7fb0c[redacted]	Billing Mode	Yearly/Monthly
Read Replica Name	dds-a7a4_readonly_node_1	Deleted Nodes	1

☑ I understand that a handling fee will be charged for this unsubscription.

📖 NOTE

You can also delete read replicas in batches by selecting all read replicas to be deleted and clicking **Batch Delete**.

- For pay-per-use instances:
  - Locate a target read replica and choose **More** > **Delete** in the **Operation** column.

Figure 7-52 Selecting a read replica in the pay-per-use instance

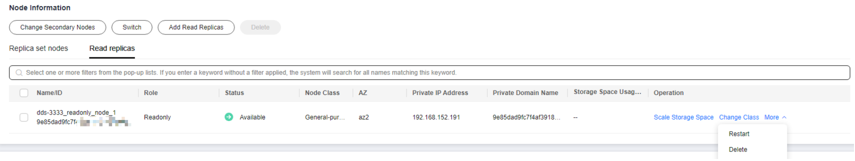

Node Information

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all names matching this keyword.

□	Name/ID	Role	Status	Node Class	AZ	Private IP Address	Private Domain Name	Storage Space Usag...	Operation
□	dds-3333_readonly_node_1 9e85dad9fc7f[redacted]	Readonly	Available	General-pur...	az2	192.168.152.191	9e85dad9fc7f4af3918...	--	Scale Storage Space Change Class More

  - Click **Yes**.

Figure 7-53 Deleting the read replica from the pay-per-use instance

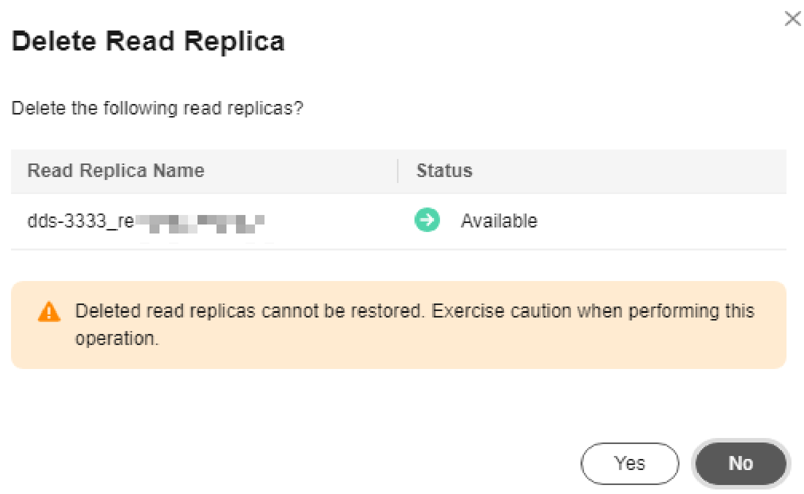

Delete Read Replica

Delete the following read replicas?

Read Replica Name	Status
dds-3333_re[redacted]	Available

Deleted read replicas cannot be restored. Exercise caution when performing this operation.

NOTE

You can also delete read replicas in batches by selecting all read replicas to be deleted and clicking **Batch Delete**.

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

----End

7.11 Configuring the Maintenance Window

During a maintenance window, Huawei Cloud O&M personnel perform maintenance operations on the instance. To prevent service interruptions, set the maintenance window to off-peak hours.

The default maintenance window is 02:00–06:00 but you can change it if needed.

Precautions

- Before maintenance is performed, DDS will send SMS and email notifications to the contact person you specified in the Huawei account.
- During the maintenance window, the DB instance may be intermittently disconnected once or twice. Ensure that your applications can reconnect to the database if the connection is interrupted.
- During DB instance maintenance, operations for service changes (such as upgrade and restart) are unavailable except account management, database management, and security group adding. When a DB instance is in maintenance, data access and query operations on the database are not affected.
- Changing the maintenance window does not affect the execution of tasks that have been scheduled.
- You can configure a maintenance window only for restarting a DB instance, changing an instance class, or upgrading the minor version of a DB instance.

- You can cancel a scheduled task to be executed on the **Task Center** page.
- The maintenance window cannot overlap the time window configured for backups. Otherwise, scheduled tasks may fail.
- Tasks delivered near the end of the maintenance window may fail to be scanned. In this case, the execution is canceled.

Changing a Maintenance Window

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. In the **DB Information** area on the **Basic Information** page, click **Change** in the **Maintenance Window** field.

Figure 7-54 Changing the maintenance window

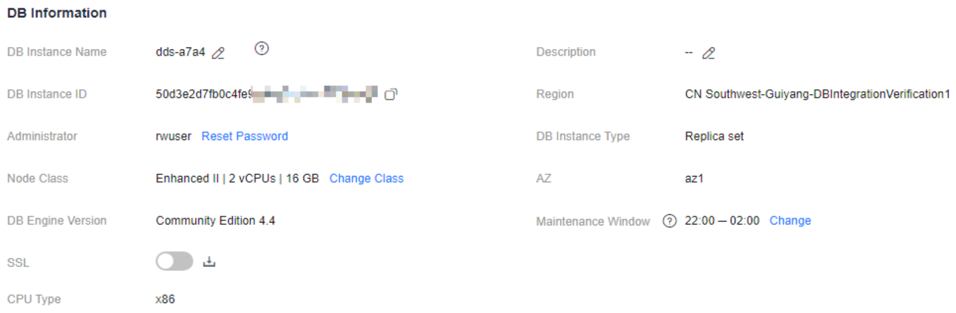

DB Information

DB Instance Name	dds-a7a4	Description	--
DB Instance ID	50d3e2d7fb0c4fe9[redacted]	Region	CN Southwest-Guiyang-DBIntegrationVerification1
Administrator	rwuser Reset Password	DB Instance Type	Replica set
Node Class	Enhanced II | 2 vCPUs | 16 GB Change Class	AZ	az1
DB Engine Version	Community Edition 4.4	Maintenance Window	22:00 – 02:00 Change
SSL	○		
CPU Type	x86		

Step 5 In the displayed dialog box, select an interval and a maintenance window, and click **OK**.

Figure 7-55 Changing the maintenance window

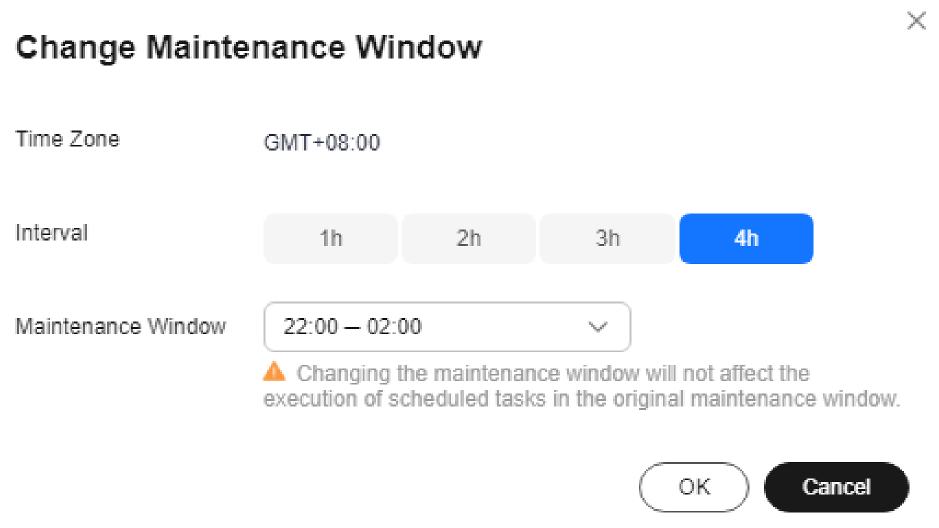

Change Maintenance Window

Time Zone GMT+08:00

Interval

Maintenance Window 22:00 – 02:00

⚠ Changing the maintenance window will not affect the execution of scheduled tasks in the original maintenance window.

----End

Canceling a Scheduled Task

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Task Center** page, locate the specified task and click **Cancel** in the **Operation** column.

Figure 7-56 Canceling a task

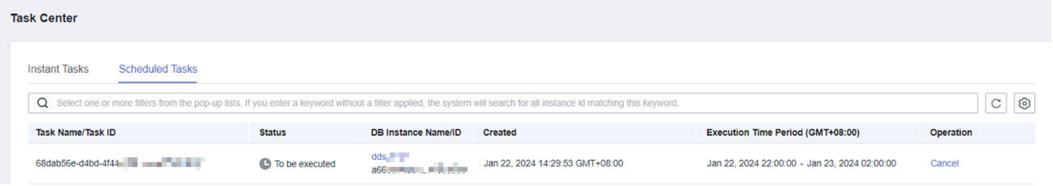

Task Center

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all instance id matching this keyword.

Task Name/Task ID	Status	DB Instance Name/ID	Created	Execution Time Period (GMT+08:00)	Operation
68dab56e-d4bd-4f44[redacted]	To be executed	dds[redacted] a66[redacted]	Jan 22, 2024 14:29:53 GMT+08:00	Jan 22, 2024 22:00:00 - Jan 23, 2024 02:00:00	Cancel

Step 5 View the result.

On the **Task Center** page, you can view the result. After the task is cancelled, its status changes to **Cancelled**.

----End

7.12 Changing an AZ

You can migrate an instance to any AZ in the same region.

Precautions

- Clusters and replica sets can be migrated between AZs.
- Instances deployed across AZs and associated with an IPv6 subnet do not support this operation.
- Inactive secondary nodes and read replicas in a replica set instance do not support this operation.
- If a cluster instance has read replicas associated, the instance cannot be migrated to another AZ.
- Services will be interrupted for up to 60 seconds while the AZ is being changed. The time required to change an AZ depends on the amount of data to be migrated. You are advised to change an AZ during off-peak hours. You are advised to use an HA connection to access the instance or configure your client to automatically reconnect to the instance.
- The destination AZ and the AZ of the current DB instance are in the same region.
- For details about regions and AZs, see **Regions and AZs**.
- To ensure stable operation of a DB instance, change an AZ during off-peak hours.

Supported Migration Types and Scenarios

Table 7-7 Supported migration types and scenarios

Migration Type	Scenario
Migrating data from one AZ to another AZ	DDS instances can be migrated to the AZ to which the ECS belongs. DDS instances and ECS in the same AZ can be connected through a private network with lower network latency.
Migrating data from a single AZ to multiple AZs	The instance disaster recovery capability needs to be improved.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the **DB Information** area on the **Basic Information** page, click **Change** to the right of the **AZ** field.

Step 6 On the displayed page, select a desired AZ and click **OK**.

Step 7 On the **Instances** page, check the changed AZ.

- During the changes, the instance status is **Changing AZ**.
- In the upper right corner of the instance list, click ↻ to refresh the list. After the migration is complete, the instance status will become **Available**.
- In the **DB Information** on the **Basic Information** page, view the new AZ where the DB instance is deployed.

----End

7.13 Updating the OS of a DB Instance

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

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

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

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

8 Data Backups

8.1 Backup Principles and Solutions

DDS instances support automated and manual backups. You can periodically back up databases. If a database is faulty or data is damaged, you can restore the database using backup files to ensure data reliability.

Backup Principles

- Cluster instance

  A cluster instance consists of a config node, and multiple dds mongos and shard nodes. The config node is used to store the configuration information of a cluster instance, and the shard node is used to store data of a cluster instance. Backing up a cluster instance means that data on the config and shard nodes is backed up separately. As shown in **Figure 8-1**, the config and shard nodes in a cluster instance are backed up to their own hidden nodes. The backup process occupies certain CPU and memory resources of the hidden nodes. During the backup, the CPU usage, memory usage, and primary/standby delay of the hidden node increase slightly, which is normal. The backup files on the hidden nodes will then be compressed and stored in OBS, and the storage space of the instance will not be occupied.

Figure 8-1 Cluster backup principle

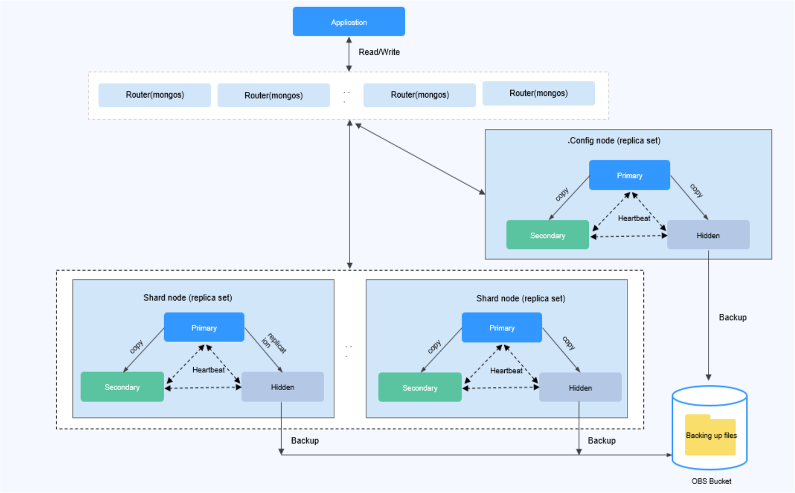

- Replica set instance

  As shown in **Figure 8-2**, replica set instance data is backed up on hidden nodes. The backup process occupies certain CPU and memory resources of the hidden node. During the backup, the CPU usage, memory usage, and primary/standby delay of the hidden node increase slightly, which is normal. The backup files on the hidden nodes will then be compressed and stored in OBS, and the storage space of the instance will not be occupied.

Figure 8-2 Replica set backup principle

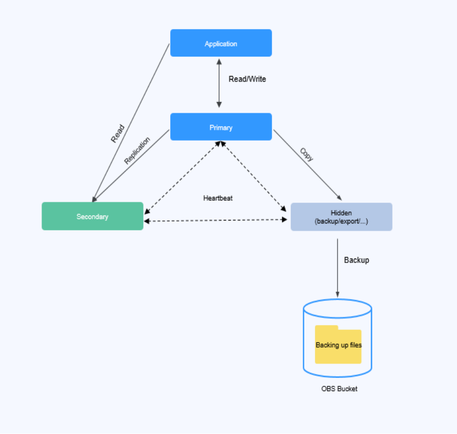

- Single node instance:

  Single-node instance backup is performed on only one node. The backup file is stored in OBS as a package, which does not occupy the storage of the instance.

NOTICE

A single node instance is backed up using mongodump. During the backup, CPU and memory resources of the node are occupied. If the resources are insufficient, the backup fails. You are advised to migrate the single-node instance data to a replica set instance for backup.

Figure 8-3 Single-node instance backup principle

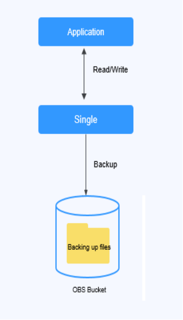

Backup and Restoration Solution

- **Table 8-1** describes how to back up and download backup files.

  NOTE

  By default, all DDS versions 3.2, 3.4, 4.0, 4.2, and 4.4 are supported unless otherwise specified.

Table 8-1 Backup solutions

Task Type	Method	Instance Type and Version	Scenario
Backing up data	**Automated backup**	- Cluster - Replica set - Single node	You can perform automated backup for DDS instances on the management console.
	Incremental backup	- Cluster (versions 4.0 and 4.2) - Replica set (versions 4.0 and 4.2)	You can perform incremental backup for DDS instances on the management console.
	Cross-Region Backup	- Cluster - Replica set - Single node	You can perform cross-region backup on the DDS console.
	Manual backup	- Cluster - Replica set - Single node	You can perform manual backup for DDS instances on the management console.
	mongodump	- Cluster - Replica set - Single node	You can use the backup and restoration tool provided by the MongoDB client to back up your self-built MongoDB database or MongoDB database on the cloud.

Task Type	Method	Instance Type and Version	Scenario
	mongoexport	- Cluster - Replica set - Single node	You can use the backup and restoration tool provided by the MongoDB client to back up your self-built MongoDB database or MongoDB database on the cloud.
Downloading a backup file	**OBS Browser+**	- Cluster - Replica set - Single node	If the size of a backup file is greater than 400 MB, use OBS Browser+ to download the file.
	Browser	- Replica set - Single node	You can directly download backup files using a browser.
	URL	- Cluster - Replica set - Single node	You can download backup files in a new browser window, or using Xunlei or Wget.

- For details about the DDS restoration scheme, see **Solutions**.

Billing

Backups are saved as packages in OBS buckets. Backups occupy backup space in OBS. If the free space DDS provides is used up, the additional space required will be billed. For the billing details, see **How Is DDS Backup Data Billed?**

8.2 Configuring an Automated Backup Policy

Scenario

DDS backs up data automatically based on the automated backup policy you set. You are advised to regularly back up data in your database. If the database becomes faulty or data is damaged, you can restore it with the backup.

The automated backup policy for DDS is enabled by default. After an instance is created, you can **modify** or **disable** the automated backup policy as required.

Once the automated backup policy is enabled, a full backup is triggered immediately. After that, full backups will be created based on the backup window and backup cycle you specify. When an instance is being backed up, data is copied and then compressed and uploaded to OBS. The length of time the backup data is kept for depends on the backup retention period you configure. The backup duration depends on the amount of data, and the average backup speed is 60 MB/s. After the automated backup policy is enabled, an incremental backup is automatically performed every 5 minutes for replica set instances to ensure data reliability. If the incremental backup function is required for cluster instances, you need to manually enable it.

Automated Backup Description

- Backup type
  - Full backup: All data is backed up even if no data is updated since the last backup.
  - Incremental backup: Incremental backup is used to back up the data newly added or modified since the last full or incremental backup. DDS automatically backs up the updated data every 5-60 minutes since the last automated or incremental backup was made.
- Backup mode
  - **Physical**: Data is copied from physical disks.
  - Snapshot
    - The data status at a particular point in time is retained. Compared with physical backup, snapshot backup is faster. After CBR is enabled, the free backup space is unavailable. You are billed for database server backup vaults on a pay-per-use basis. For details, see **How Is CBR Billed?**

      NOTE

      - The backup time is proportional to how much data your instance has. Too much data can decrease the backup efficiency. If you have large amounts of data and want to speed up the backup process, contact customer service to enable Cloud Backup and Recovery (CBR).
      - After CBR is enabled, snapshot backup is used. Existing automated and manual backups can still be used to restore data.
      - When you delete a DB instance, its automated backups are also deleted but its manual backups are retained.
      - After CBR is enabled, the next full backup is a snapshot backup. You can use the snapshot backup to restore data.
    - If more than 2 TB of data needs to be backed up, the backup method cannot be set back to physical backup.
    - Snapshots cannot be backed up across regions.
  - **Logical**: A tool is used to read data and logically export the data.
- **Table 8-2** lists the automated backup methods supported by DDS.

Table 8-2 Backup methods

Instance Type	Backup Mode	Backup Type
Cluster	Physical backup/ Snapshot backup **NOTE** - Only whitelisted users can use snapshot backup. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket. - For details about how to set **Backup Method** to **Snapshot**, see **Setting Backup Method for a DB Instance**.	- Full backup - Incremental backup **NOTE** Snapshot backup only applies to physical data in a full backup.
Replica set	Physical backup/ Snapshot backup **NOTE** - Only whitelisted users can use snapshot backup. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket. - For details about how to set **Backup Method** to **Snapshot**, see **Setting Backup Method for a DB Instance**.	- Full backup - Incremental backup **NOTE** Snapshot backup only applies to physical data in a full backup.

Instance Type	Backup Mode	Backup Type
Single node **NOTE** Single node instances apply to only a few scenarios. You are advised to use a single node instance only for learning.	Logical backup/ Snapshot backup **NOTE** • Only whitelisted users can use snapshot backup. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket. • For details about how to set **Backup Method** to **Snapshot**, see **Setting Backup Method for a DB Instance**.	Full backup

Pricing

- After you purchase an instance, DDS will provide additional backup storage of the same size as you purchased. For example, if you purchase 100 GB of instance storage space, you will obtain 100 GB of backup storage space. If the size of backup data does not exceed 100 GB, the backup data is stored on OBS free of charge. If the size of the backup data exceeds 100 GB, you will be charged based on the **OBS billing rules**.
- You can check your expenditure records for DDS backup fees by going to **Billing Center** > **Bills**.

Precautions

- The backup process does not affect services.
- DDS checks existing automated backup files. If the retention period of a file exceeds the backup retention period you set, DDS will delete the file.
- After the backup policy is modified, an automated backup will be triggered based on the new backup policy. The retention period of the previously generated automated backups remains unchanged.
- Single node instances do not support incremental backup.
- By default, the name of an automated backup ends with the UTC time. To change the display time in the automated backup name to the local time, contact Huawei O&M personnel.

Enabling or Modifying an Automated Backup Policy

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Set Backup Policy**. If you want to enable the automated backup policy, click ○. Once enabled, the backup policy can be modified as shown in **Figure 8-4**.

Figure 8-4 Set Backup Policy

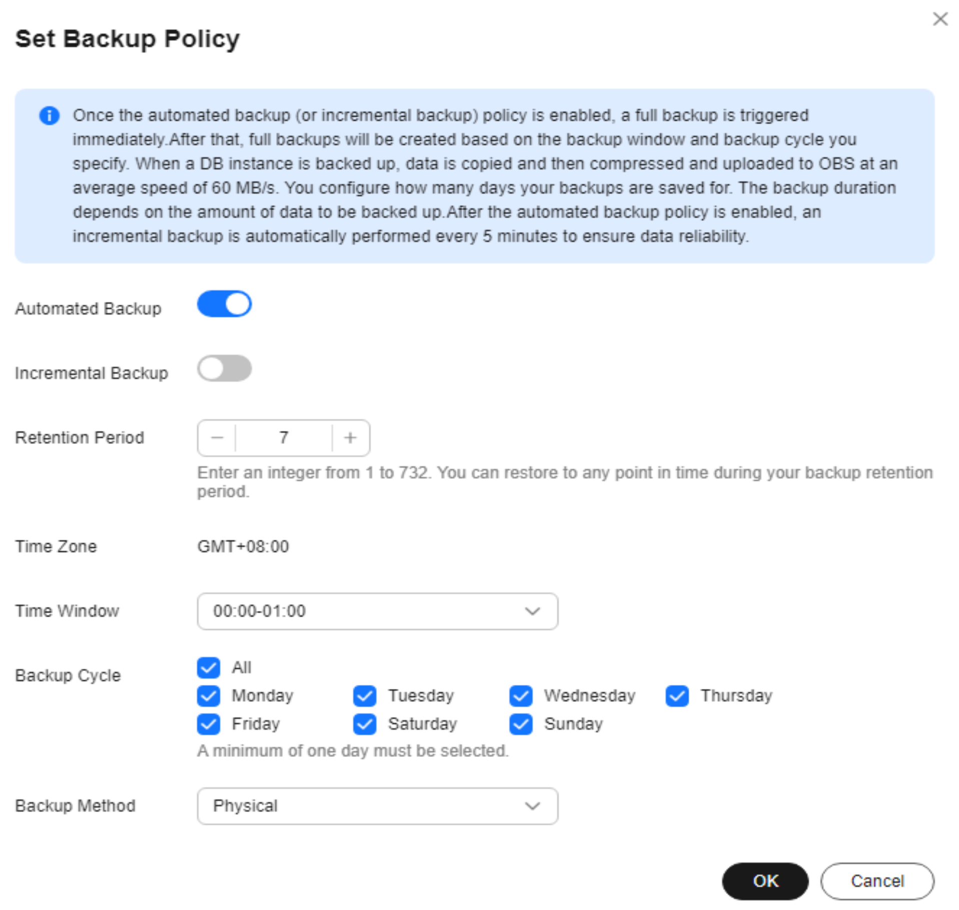

Set Backup Policy

Once the automated backup (or incremental backup) policy is enabled, a full backup is triggered immediately.After that, full backups will be created based on the backup window and backup cycle you specify. When a DB instance is backed up, data is copied and then compressed and uploaded to OBS at an average speed of 60 MB/s. You configure how many days your backups are saved for. The backup duration depends on the amount of data to be backed up.After the automated backup policy is enabled, an incremental backup is automatically performed every 5 minutes to ensure data reliability.

Automated Backup: ●

Incremental Backup: ○

Retention Period: 7

Enter an integer from 1 to 732. You can restore to any point in time during your backup retention period.

Time Zone: GMT+08:00

Time Window: 00:00-01:00

Backup Cycle: ☑ All
☑ Monday ☑ Tuesday ☑ Wednesday ☑ Thursday
☑ Friday ☑ Saturday ☑ Sunday

A minimum of one day must be selected.

Backup Method: Physical

Table 8-3 Parameter description

Parameter	Description
Retention Period (days)	Number of days that your automated backups can be retained. The retention period is from 1 to 732 days and the default value is **7**. ● Extending the retention period improves data reliability. You can configure the retention period if needed. ● If you shorten the retention period, the new backup policy takes effect for existing backups. Any automated backups (including full and incremental backups) that have expired will be automatically deleted. Manual backups will not be automatically deleted but you can delete them manually.
Time Zone	The default backup time zone is the UTC time.
Time Window	A one-hour period the backup will be scheduled within 24 hours, such as 01:00-02:00. The backup time is in UTC format.
Backup Cycle	● If you set the retention period to 1 to 6 days, data is automatically backed up each day of the week and the backup cycle cannot be changed. ● If you set the retention period to 7 to 732 days, you must select at least one day of the week for the backup cycle.
Backup Method	● **Physical**: Data is copied from physical disks. ● **Snapshot**: The data status at a particular point in time is retained. Compared with physical backup, snapshot backup is faster. ● **Logical**: A tool is used to read data and logically export the data.

Policy for automatically deleting full backups:

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

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

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

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

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

Step 7 Click **OK** to save the changes.

Step 8 View the results.

- During the creation of an automated backup, you can query the backup status on the **Backups** page or the **Backups & Restorations** tab. The backup status is **Backing up**.
- In the upper right corner of the backup list, click ↻ to refresh the list. The backup status changes to **Complete**. The backup type is **Automated** and the backup method is **Physical**.

----End

Disabling an Automated Backup Policy

NOTICE

When disabling the automated backup policy:

- Your data cannot be backed up.
- Your replica set instances cannot be restored to a specified point in time.
- If you choose to delete all the existing automated backup when disabling the automated backup policy, related restoration or download operations will fail.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Set Backup Policy**. On the displayed page, click ⬤ to disable the automated backup policy. **Figure 8-5** shows the dialog box for modifying the backup policy.

Figure 8-5 Set Backup Policy

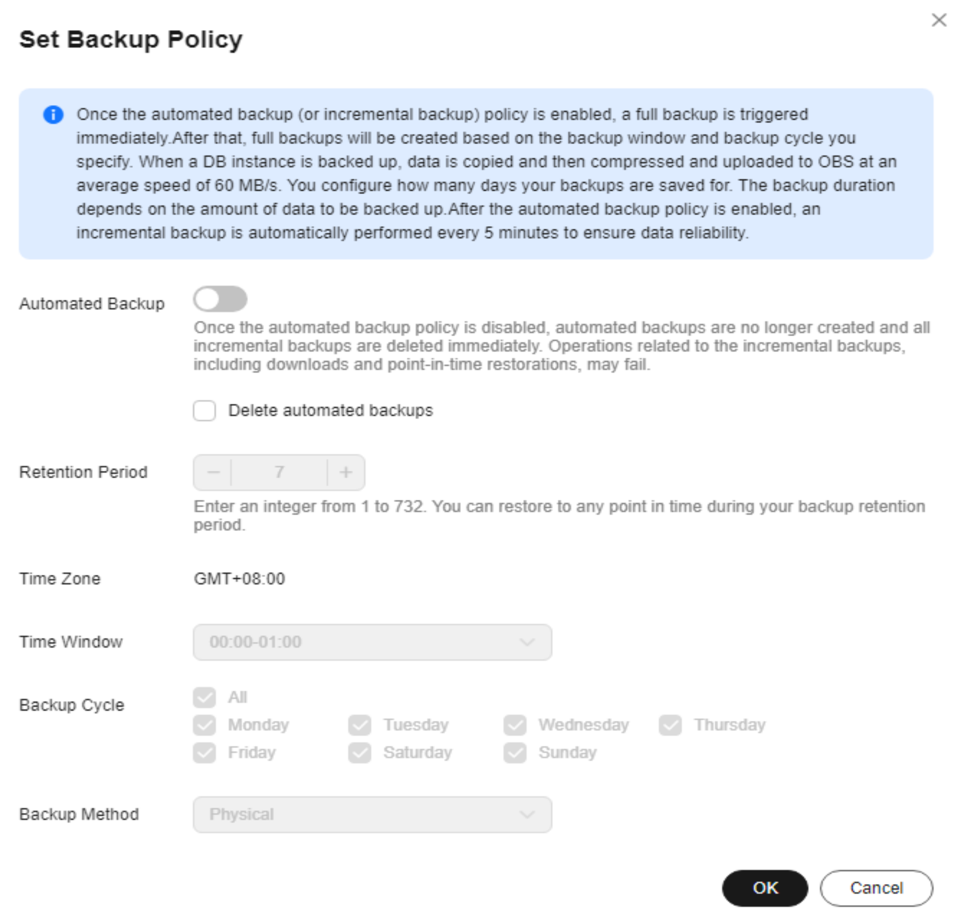

Set Backup Policy

Once the automated backup (or incremental backup) policy is enabled, a full backup is triggered immediately.After that, full backups will be created based on the backup window and backup cycle you specify. When a DB instance is backed up, data is copied and then compressed and uploaded to OBS at an average speed of 60 MB/s. You configure how many days your backups are saved for. The backup duration depends on the amount of data to be backed up.After the automated backup policy is enabled, an incremental backup is automatically performed every 5 minutes to ensure data reliability.

Automated Backup

Once the automated backup policy is disabled, automated backups are no longer created and all incremental backups are deleted immediately. Operations related to the incremental backups, including downloads and point-in-time restorations, may fail.

□ Delete automated backups

Retention Period: 7

Enter an integer from 1 to 732. You can restore to any point in time during your backup retention period.

Time Zone: GMT+08:00

Time Window: 00:00-01:00

Backup Cycle: ☑ All ☑ Monday ☑ Tuesday ☑ Wednesday ☑ Thursday ☑ Friday ☑ Saturday ☑ Sunday

Backup Method: Physical

You can determine whether to delete all automated backup files:

- If you do not select **Delete automated backups**, all backup files within the retention period will be retained, but you can still delete them manually. For details, see section **Deleting an Automated Backup**.
- If you select **Delete automated backups**, all backup files within the retention period will be deleted.

If you shorten the retention period, the new backup policy takes effect for all backup files. Any backup files that have expired, based on a newly configured retention period, will be deleted, but the latest expired backup file will be retained.

Step 7 Click **OK**.

NOTE

- If automated backups are disabled, any automated backups in progress stop immediately.
- After automated backups are disabled, incremental backups are disabled by default.
- If you need to enable the automated backup policy again, see **Enabling or Modifying an Automated Backup Policy**.

----End

8.3 Configuring an Incremental Backup Policy

Incremental backup is used to back up the data newly added or modified since the last full or incremental backup. DDS automatically backs up the updated data every 5-60 minutes since the last automated or incremental backup was made.

When you create a DDS DB instance, incremental backup is enabled by default for all DB instances except DB instances with fewer than 4 vCPUs. You can enable or disable the backup policy after an instance is created. For details, see **Enabling or Modifying an Incremental Backup Policy** and **Disabling the Incremental Backup Policy**.

After an incremental backup policy is enabled for a DDS instance, incremental files are not displayed on the DDS console.

Prerequisites

Before enabling the incremental backup policy, ensure that the automated backup policy has been enabled. For details, see **Enabling or Modifying an Automated Backup Policy**.

NOTE

Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Precautions

- Incremental backup is performed on the hidden node. After an incremental backup policy is enabled, the CPU and memory usage of the hidden node increases, depending on the service model.
- To ensure stable database running, you are advised to double the specifications of the hidden node when the CPU usage exceeds 60% or the memory usage exceeds 80%.

Constraints

- Only cluster instances of version 4.0 or later and replica set instances of version 3.4 or later support this function.
- To minimize the impact of incremental backup on instances, incremental backup is disabled by default for DB instances with fewer than 4 vCPUs.
- Incremental backup stops in any of the following scenarios and starts again after the next automated backup is complete:
  - rename operation
  - collmod operation
  - Creating a user
  - Deleting a user
  - Creating a role
  - Deleting a role

- Enabling shard IP addresses of a cluster instance
- Changing the password of the shard node user
- Enabling config IP addresses of a cluster instance
- Changing the password of the config node user
- Changing the password of the **rwuser** user

Enabling or Modifying an Incremental Backup Policy

Step 1 Log in to the management console.

Step 2 Click ⊙ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Set Backup Policy**. To enable incremental backup, click ○. After incremental backup is enabled, a full backup is triggered.

Figure 8-6 Set Backup Policy

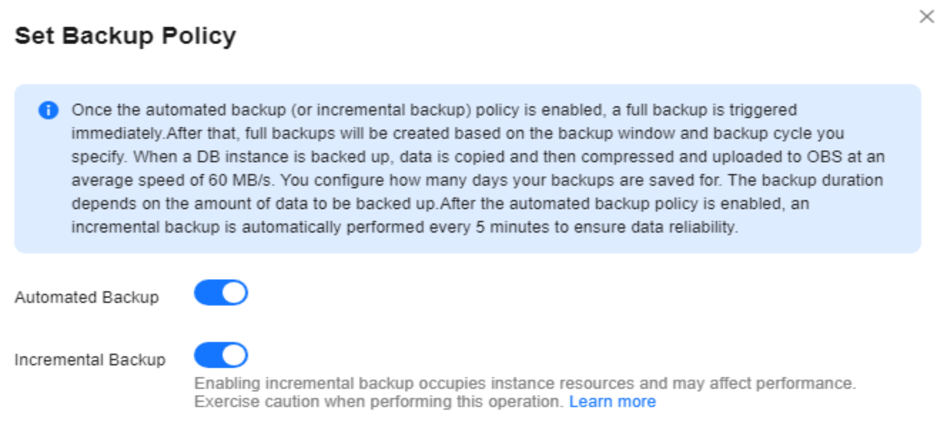

Set Backup Policy

Once the automated backup (or incremental backup) policy is enabled, a full backup is triggered immediately.After that, full backups will be created based on the backup window and backup cycle you specify. When a DB instance is backed up, data is copied and then compressed and uploaded to OBS at an average speed of 60 MB/s. You configure how many days your backups are saved for. The backup duration depends on the amount of data to be backed up.After the automated backup policy is enabled, an incremental backup is automatically performed every 5 minutes to ensure data reliability.

Automated Backup ●

Incremental Backup ●

Enabling incremental backup occupies instance resources and may affect performance. Exercise caution when performing this operation. Learn more

Table 8-4 Parameter description

Parameter	Description
Automated Backup	For details about automated backup parameters, see Table 8-3.
Incremental Backup	Before enabling the incremental backup policy, ensure that the automated backup policy has been enabled.

Step 7 Click **OK**.

Step 8 View the results.

- During the creation of an automated backup, you can query the backup status on the **Backups** page or the **Backups & Restorations** tab. The backup status is **Backing up**.
- In the upper right corner of the backup list, click ↻ to refresh the list. The backup status changes to **Complete**.

----End

Disabling the Incremental Backup Policy

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Set Backup Policy**.

Step 7 In the displayed dialog box, click 🔘 to the right of **Incremental Backup** to disable the incremental backup policy.

Figure 8-7 Disabling incremental backup policy

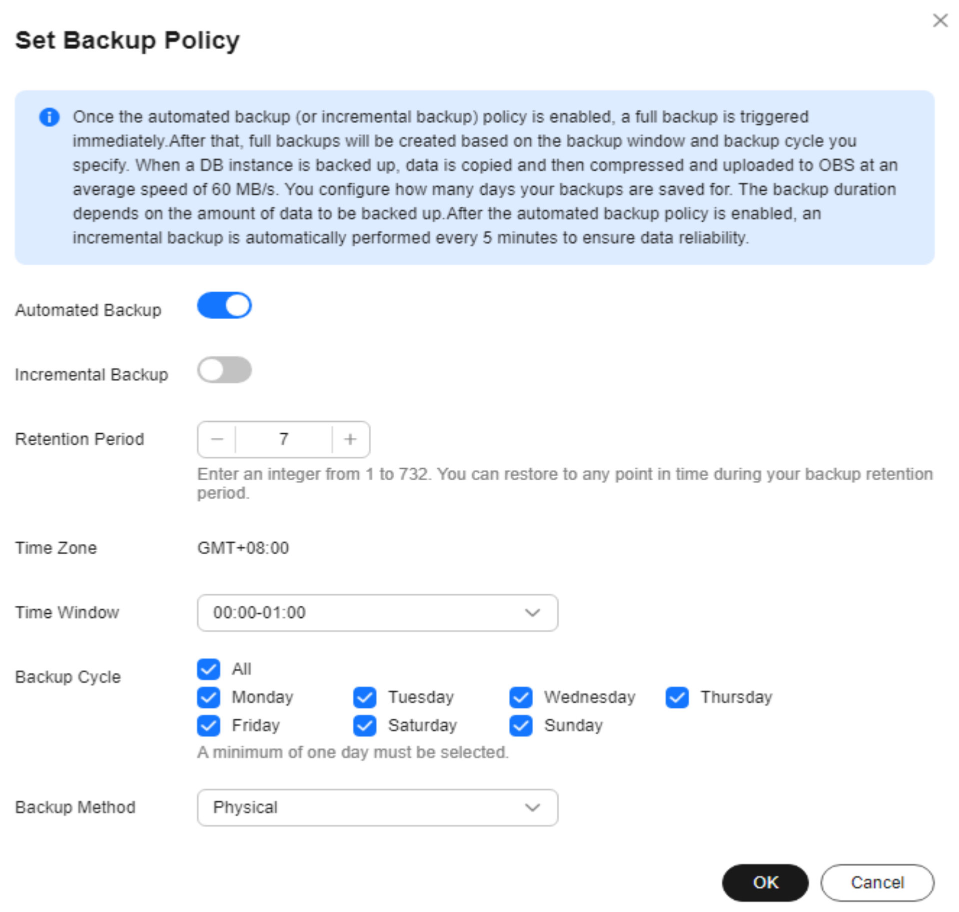

Set Backup Policy

Once the automated backup (or incremental backup) policy is enabled, a full backup is triggered immediately.After that, full backups will be created based on the backup window and backup cycle you specify. When a DB instance is backed up, data is copied and then compressed and uploaded to OBS at an average speed of 60 MB/s. You configure how many days your backups are saved for. The backup duration depends on the amount of data to be backed up.After the automated backup policy is enabled, an incremental backup is automatically performed every 5 minutes to ensure data reliability.

Automated Backup: ●

Incremental Backup: ○

Retention Period: − 7 +

Enter an integer from 1 to 732. You can restore to any point in time during your backup retention period.

Time Zone: GMT+08:00

Time Window: 00:00-01:00

Backup Cycle: ☑ All
☑ Monday ☑ Tuesday ☑ Wednesday ☑ Thursday
☑ Friday ☑ Saturday ☑ Sunday

A minimum of one day must be selected.

Backup Method: Physical

Step 8 Click **OK**.

NOTICE

- After you disable this incremental backup, the incremental backup task will be stopped, all incremental backup files will be deleted immediately, and operations related to incremental backup fail.
- After a DB instance is deleted, all incremental backup files of the DB instance are retained. The retention period depends on the incremental backup retention period that you specified.

----End

8.4 Configuring the Cross-Region Backup Policy

DDS can store backup files in the destination region or OBS, so you can use the backup files in the destination region to restore data to a new DDS instance.

After the cross-region backup policy is enabled, the system automatically stores the backup files created for the instance to the destination region you specified. You can manage cross-region backup files on the **Backups** page.

Before You Start

- To apply for the permission to set cross-region backup policies, contact customer service.
- Before enabling the cross-region backup policy, ensure that the automated backup policy has been enabled. Otherwise, the cross-region backup cannot take effect. For details, see **Enabling or Modifying an Automated Backup Policy**.

Billing

Table 8-5 Billing

Specification Code	Billing Item	Unit Price
dds.mongodb.crossreg.backup.space.repset	Storage space	For details, see **Product Pricing Details**.
dds.mongodb.crossreg.backup.space.single	Storage space	
dds.mongodb.crossreg.backup.space	Storage space	

Enabling or Modifying a Cross-Region Backup Policy

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the target instance.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Set Cross-Region Backup Policy**.

Figure 8-8 Set Cross-Region Backup Policy

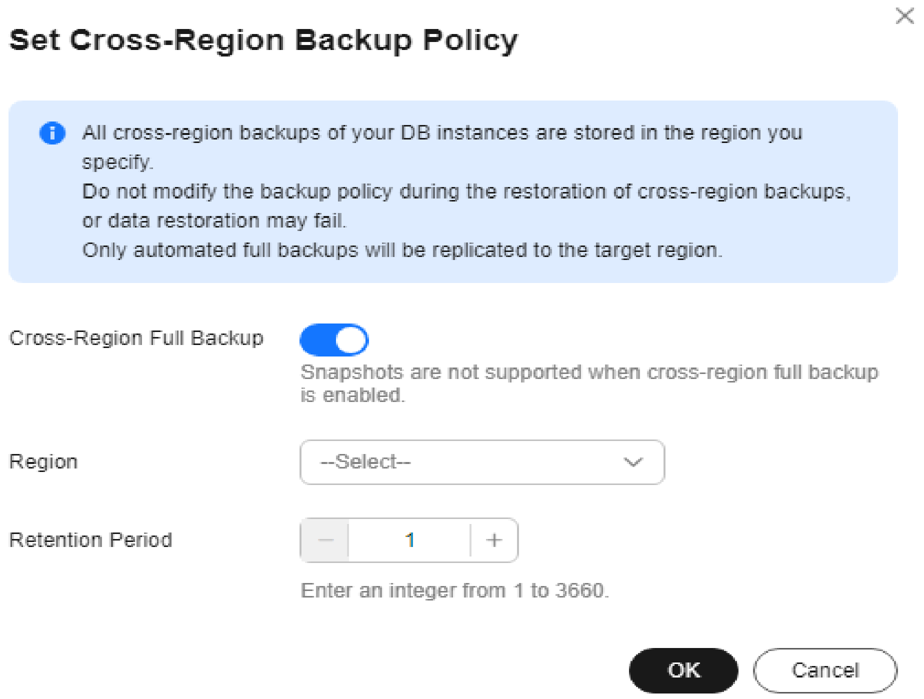

Set Cross-Region Backup Policy

All cross-region backups of your DB instances are stored in the region you specify.
Do not modify the backup policy during the restoration of cross-region backups, or data restoration may fail.
Only automated full backups will be replicated to the target region.

Cross-Region Full Backup
Snapshots are not supported when cross-region full backup is enabled.

Region: --Select--

Retention Period: 1
Enter an integer from 1 to 3660.

Table 8-6 Parameter description

Parameter	Description
Cross-Region Full Backup	Click ○ to back up the automated full backup file of the instance to a remote location.
Cross-Region Incremental Backup	Click ○ to back up the incremental backup file of the instance to a remote location. **NOTE** ● Only replica set instances support cross-region incremental backup. ● If cross-region full backup is not enabled, cross-region incremental backup cannot be enabled. ● After cross-region incremental backup is enabled, you can restore an instance to a specified time point only after the next automated full backup replication is complete. The specified time point must be later than the time when the automated full backup is complete.
Region	Select the region for which you back up data based on service requirements.
Retention Period	**Retention Period** refers to the number of days (range: 1 to 3,660) that data is kept. You can increase the retention period to improve data reliability.

Step 7 Click **OK**.

Step 8 On the **Cross-Region Backups** tab of the **Backups** page, manage cross-region backup files.

- To modify the cross-region backup policy, click **Set Cross-Region Backup** in the **Operation** column.
- To view generated cross-region backup files, click **View Cross-Region Backup** in the **Operation** column. You can use the cross-region backup files to restore data to a new instance.

----End

Disabling a Cross-Region Backup Policy

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the target instance.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Set Cross-Region Backup Policy**.

Step 7 In the displayed dialog box, click to disable the cross-region backup policy.

Figure 8-9 Disabling a cross-region backup policy

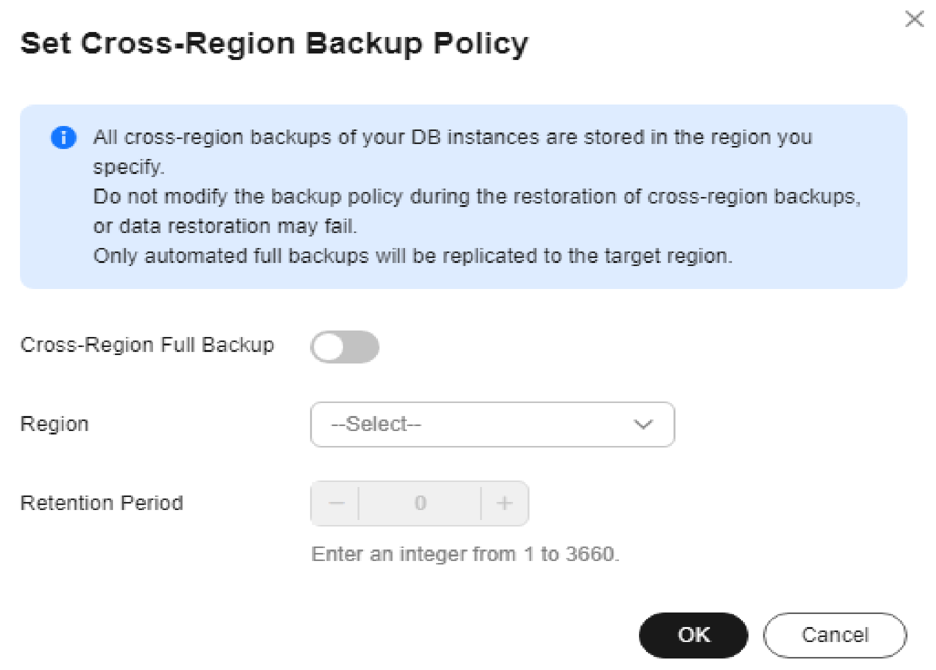

Set Cross-Region Backup Policy

All cross-region backups of your DB instances are stored in the region you specify.
Do not modify the backup policy during the restoration of cross-region backups, or data restoration may fail.
Only automated full backups will be replicated to the target region.

Cross-Region Full Backup

Region: --Select--

Retention Period: 0

Enter an integer from 1 to 3660.

Step 8 Click **OK**.

NOTICE

- If the cross-region backup policy is disabled, the cross-region backup task will be stopped immediately, and all cross-region backup and cross-region incremental backup files will be immediately deleted. Operations related to cross-region backup or incremental backup may fail.
- After an instance is deleted, all cross-region backups and incremental backups of the instance will be retained. The retention period is determined by the retention period you specified in the cross-region backup policy.

----End

8.5 Setting Backup Method for a DB Instance

DDS allows snapshot backup for a DB instance.

NOTE

Huawei Cloud has discontinued the sale of DDS single node instances since July 15, 2023.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Set Backup Policy**.

Figure 8-10 Setting backup method for a DB instance

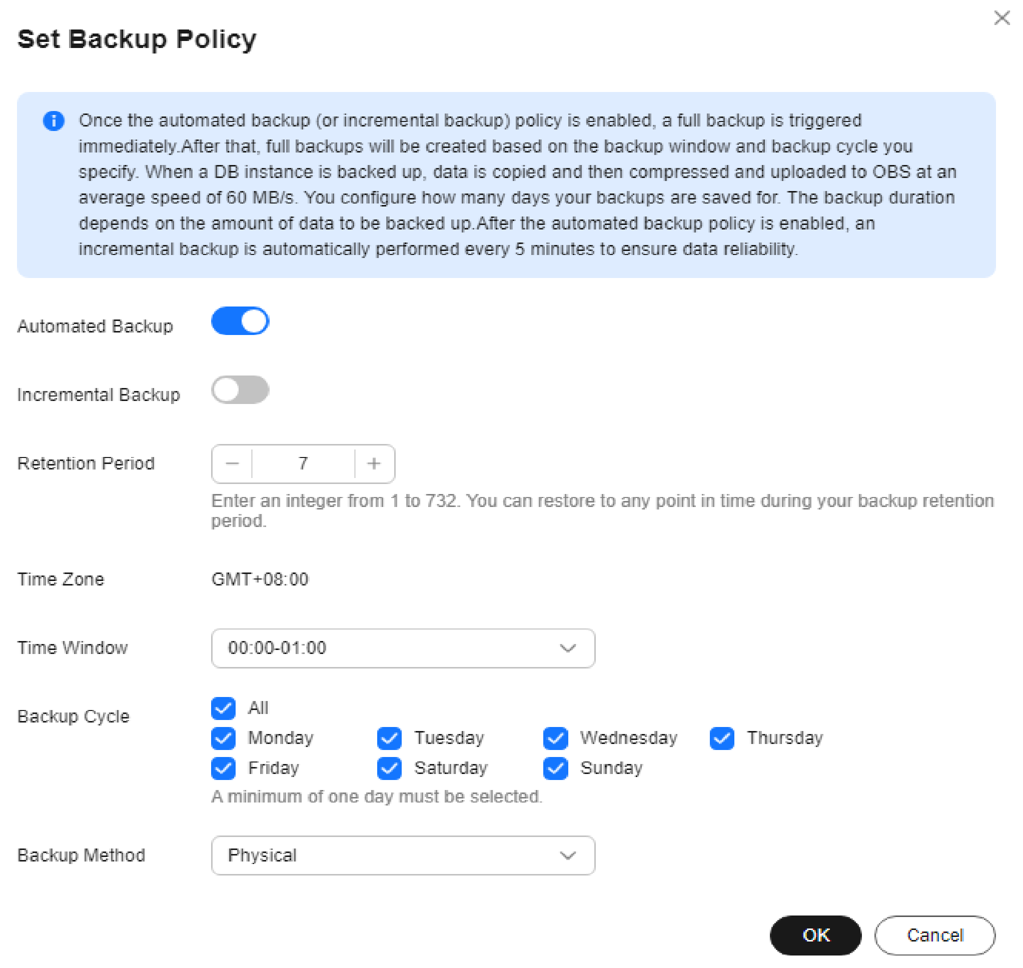

Set Backup Policy

Once the automated backup (or incremental backup) policy is enabled, a full backup is triggered immediately.After that, full backups will be created based on the backup window and backup cycle you specify. When a DB instance is backed up, data is copied and then compressed and uploaded to OBS at an average speed of 60 MB/s. You configure how many days your backups are saved for. The backup duration depends on the amount of data to be backed up.After the automated backup policy is enabled, an incremental backup is automatically performed every 5 minutes to ensure data reliability.

Automated Backup: On

Incremental Backup: Off

Retention Period: 7

Enter an integer from 1 to 732. You can restore to any point in time during your backup retention period.

Time Zone: GMT+08:00

Time Window: 00:00-01:00

Backup Cycle: ☑ All
☑ Monday ☑ Tuesday ☑ Wednesday ☑ Thursday
☑ Friday ☑ Saturday ☑ Sunday

A minimum of one day must be selected.

Backup Method: Physical

Table 8-7 Parameter description

Parameter	Description
Retention Period (days)	Number of days that your automated backups can be retained. The retention period is from 1 to 732 days and the default value is **7**. ● Extending the retention period improves data reliability. You can configure the retention period if needed. ● If you shorten the retention period, the new backup policy takes effect for existing backups. Any automated backups (including full and incremental backups) that have expired will be automatically deleted. Manual backups will not be automatically deleted but you can delete them manually.
Time Zone	The default backup time zone is the UTC time.
Time Window	The backup interval is one hour. You are advised to set the backup window to an off-peak period.

Parameter	Description
Backup Cycle	● If you set the retention period to 1 to 6 days, data is automatically backed up each day of the week and the backup cycle cannot be changed. ● If you set the retention period to 7 to 732 days, you must select at least one day of the week for the backup cycle.
Backup Method	● **Physical**: Data is copied from physical disks. ● **Snapshot**: The data status at a particular point in time is retained. Compared with physical backup, snapshot backup is faster. ● **Logical**: A tool is used to read data and logically export the data.

Step 7 Set **Backup Method** to **Snapshot** and click **OK**.

----End

8.6 Creating a Manual Backup

This section describes how to create a manual backup. Creating a backup for a DB instance helps ensure data can be restored if needed, ensuring data reliability.

Prerequisites

You can create backups (including manual backups, automated backups, and incremental backups) only when the hidden nodes of cluster instances and replica set instances are normal.

Description

- Backup type

  Full backup: All data is backed up even if no data is updated since the last backup.

- Backup mode

  Physical backup: Data is copied from physical disks.

- Table 8-8 lists the manual backup methods supported by DDS.

Table 8-8 Backup methods

Instance Type	Backup Mode	Backup Type
Cluster	Physical backup	Full backup
Replica set	Physical backup	Full backup

Instance Type	Backup Mode	Backup Type
Single node **NOTE** Single node instances apply to only a few scenarios. You are advised to use a single node instance only for learning.	Logical backup	Full backup

Pricing Details

- After you purchase an instance, DDS will provide additional backup storage of the same size as you purchased. For example, if you purchase 100 GB of instance storage space, you will obtain 100 GB of backup storage space. If the size of backup data does not exceed 100 GB, the backup data is stored on OBS free of charge. If the size of the backup data exceeds 100 GB, you will be charged based on the **OBS billing rules**.
- You can check your expenditure records for DDS backup fees by going to **Billing Center** > **Bills**.
- Backups that are not normally delivered by a customer (for example, full backups automatically delivered after node rebuilding) are not displayed on the **Backups** page because they are not billed.

Precautions

- The backup process does not affect services.
- When you delete a DB instance, its automated backups are also deleted but its manual backups are retained.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 Create a manual backup on the DDS console in any of the following ways:

- On the **Instances** page, locate an available instance and choose **More** > **Create Backup** in the **Operation** column.

Figure 8-11 Method 1: Creating a backup

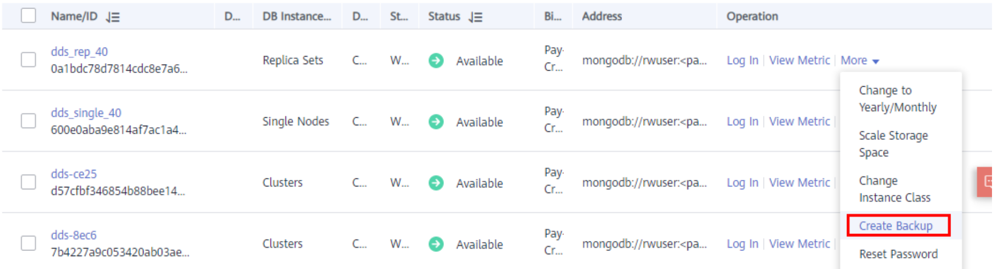

□	Name/ID	D...	DB Instance...	D...	St...	Status	Bi...	Address	Operation
□	dds_rep_40 0a1bdc78d7814cdc8e7a6...		Replica Sets	C...	W...	Available	Pay-Cr...	mongodb://rwuser:<pa...	Log In | View Metric | More
□	dds_single_40 600e0aba9e814af7ac1a4...		Single Nodes	C...	W...	Available	Pay-Cr...	mongodb://rwuser:<pa...	Log In | View Metric
□	dds-ce25 d57cfbf346854b88bee14...		Clusters	C...	W...	Available	Pay-Cr...	mongodb://rwuser:<pa...	Log In | View Metric
□	dds-8ec6 7b4227a9c053420ab03ae...		Clusters	C...	W...	Available	Pay-Cr...	mongodb://rwuser:<pa...	Log In | View Metric

- On the **Instances** page, choose **Backups** in the navigation pane on the left. On the displayed page, click **Create Backup**.

Figure 8-12 Method 2: Creating a backup

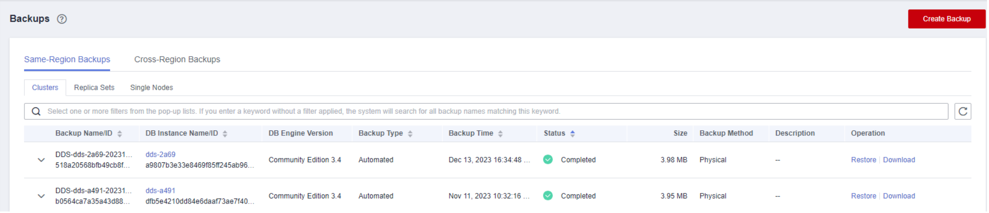
Backups
Create Backup
Same-Region Backups
Cross-Region Backups
Clusters
Replica Sets
Single Nodes
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all backup names matching this keyword.

Backup Name/ID	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dds-2a69-20231... 518a20568bfb49cb8f...	dds-2a69 a9807b3e33e8469f85ff245ab96...	Community Edition 3.4	Automated	Dec 13, 2023 16:34:48 ...	Completed	3.98 MB	Physical	--	Restore | Download
DDS-dds-a491-20231... b0564ca7a35a43d88...	dds-a491 dfb5e4210dd84e6daaf73ae7f40...	Community Edition 3.4	Automated	Nov 11, 2023 10:32:16 ...	Completed	3.95 MB	Physical	--	Restore | Download

- On the **Instances** page, click an available DB instance. In the navigation pane on the left, choose **Backups & Restorations**. On the **Backups & Restorations** page, click **Create Backup**.

Figure 8-13 Method 3: Creating a backup

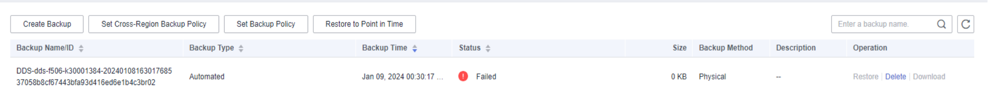

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dds-f506-k30001384-20240108163017685 37058b8cf67443bfa93d416ed6e1b4c3br02	Automated	Jan 09, 2024 00:30:17 ...	Failed	0 KB	Physical	--	Restore | Delete | Download

Step 5 In the displayed dialog box, specify **Backup Name** and **Description** and click **OK**.

- The manual backup name can be 4 to 64 characters long. It must start with a letter and can contain only letters, digits, hyphens (-), and underscores (_).
- The description contains a maximum of 256 characters and cannot contain the carriage return character and the following special characters: >!<"&'=

Step 6 View the results.

- During the creation of a manual backup, you can query the backup status on the **Backups** or the **Backups & Restorations** page. The backup status is **Backing up**. The time it takes to complete the backup depends on the size of the job.
- If the manual backup is successfully created, the backup status is **Complete**. The manual backup type is **Manual** and the backup method is **Physical**.

----End

8.7 Deleting a Manual Backup

This section describes how to delete manual backups to release the storage space.

Precautions

- Deleted backups cannot be restored. Exercise caution when performing this operation.
- Backups being used to recover instances cannot be deleted.
- To delete manual backups in batches, submit an application by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the console.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 Delete a manual backup.

On the DDS console, you can delete a manual backup using any of the following methods:

- Method 1:
  a. In the navigation pane on the left, choose **Backups**.
  b. On the **Backups** page, click the **Clusters**, **Replica Sets**, or **Single Nodes** tab.
  c. Locate the manual backup to be deleted and click **Delete** in the **Operation** column.

Figure 8-14 Deleting a Manual Backup

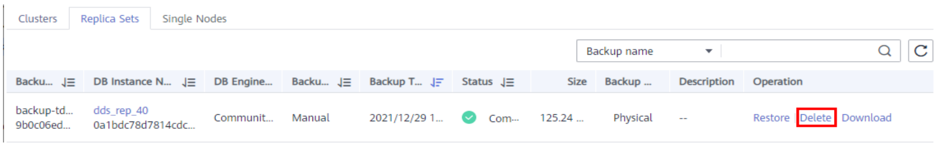

Backu...	DB Instance N...	DB Engine...	Backu...	Backup T...	Status	Size	Backup ...	Description	Operation
backup-td... 9b0c06ed...	dds_rep_40 0a1bdc78d7814cdc...	Communit...	Manual	2021/12/29 1...	✓ Com...	125.24 ...	Physical	--	Restore Delete Download

- Method 2:
  a. On the **Instances** page, click the target DB instance.
  b. In the navigation pane on the left, choose **Backups & Restorations**.
  c. On the **Backups & Restorations** page, locate the manual backup to be deleted and click **Delete**.

Figure 8-15 Deleting a Manual Backup

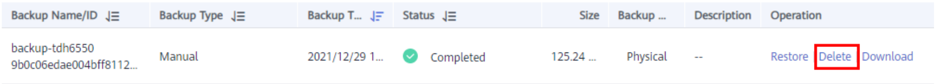

Backup Name/ID	Backup Type	Backup T...	Status	Size	Backup ...	Description	Operation
backup-tdh6550 9b0c06edae004bff8112...	Manual	2021/12/29 1...	✓ Completed	125.24 ...	Physical	--	Restore Delete Download

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

----End

8.8 Deleting an Automated Backup

DDS allows you to delete failed automated backups to release storage space. Deleted backups cannot be restored. Exercise caution when performing this operation.

Method 1

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, locate the automated backup to be deleted and click **Delete**.

Figure 8-16 Deleting an automated backup

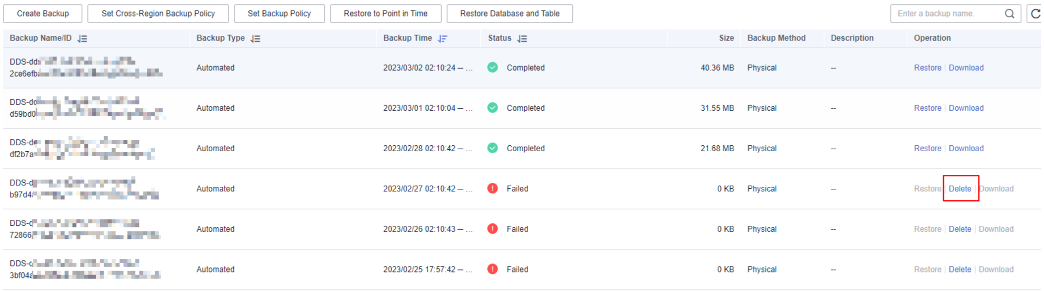

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dd[redacted] 2ce6efb[redacted]	Automated	2023/03/02 02:10:24 – ...	Completed	40.36 MB	Physical	--	Restore | Download
DDS-d[redacted] d59bd0[redacted]	Automated	2023/03/01 02:10:04 – ...	Completed	31.55 MB	Physical	--	Restore | Download
DDS-d[redacted] df2b7a[redacted]	Automated	2023/02/28 02:10:42 – ...	Completed	21.68 MB	Physical	--	Restore | Download
DDS-d[redacted] b97d4[redacted]	Automated	2023/02/27 02:10:42 – ...	Failed	0 KB	Physical	--	Restore | Delete | Download
DDS-d[redacted] 72866[redacted]	Automated	2023/02/26 02:10:43 – ...	Failed	0 KB	Physical	--	Restore | Delete | Download
DDS-d[redacted] 3bf04[redacted]	Automated	2023/02/25 17:57:42 – ...	Failed	0 KB	Physical	--	Restore | Delete | Download

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

----End

Method 2

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left of the **Instances** page, choose **Backups**.

Step 5 On the **Backups** page, click the **Clusters**, **Replica Sets**, or **Single Nodes** tab.

Step 6 Locate the automated backup to be deleted and click **Delete** in the **Operation** column.

Figure 8-17 Deleting an automated backup

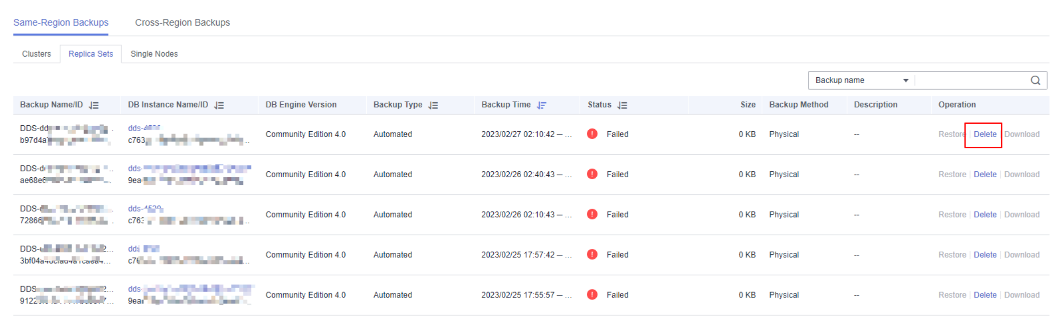

Same-Region Backups | Cross-Region Backups

Clusters | Replica Sets | Single Nodes

Backup name

Backup Name/ID	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dd[redacted] b97d4a[redacted]	dds-[redacted] c763[redacted]	Community Edition 4.0	Automated	2023/02/27 02:10:42 – ...	Failed	0 KB	Physical	--	Restore | Delete | Download
DDS-d[redacted] ae68e[redacted]	dds-[redacted] 9ea[redacted]	Community Edition 4.0	Automated	2023/02/26 02:40:43 – ...	Failed	0 KB	Physical	--	Restore | Delete | Download
DDS-[redacted] 72866[redacted]	dds-[redacted] c76[redacted]	Community Edition 4.0	Automated	2023/02/26 02:10:43 – ...	Failed	0 KB	Physical	--	Restore | Delete | Download
DDS-[redacted] 3bf04[redacted]	dds[redacted] c76[redacted]	Community Edition 4.0	Automated	2023/02/25 17:57:42 – ...	Failed	0 KB	Physical	--	Restore | Delete | Download
DDS[redacted] 9122[redacted]	dds-[redacted] 9ea[redacted]	Community Edition 4.0	Automated	2023/02/25 17:55:57 – ...	Failed	0 KB	Physical	--	Restore | Delete | Download

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

----End

8.9 Stopping a Backup

Scenarios

DDS allows you to stop a backup. If an emergency operation, such as specification change or minor version upgrade, cannot be performed because the DB instance is backing up, you can stop the backup.

Precautions

- Only full backups (streaming backups) can be stopped.
- Cross-region backups cannot be stopped.
- Only backups in the **Backing up** or **Uploading** state can be stopped.
- Stopping a backup may stop incremental backup at the current point in time and may fail. Exercise caution when performing this operation.
- Stopping a backup makes a DB instance return to the **Available** state as soon as possible to prevent blocking the execution of other tasks. The backup task may not be terminated.
- You are not allowed to stop a critical backup. If you do need, contact O&M personnel.
- To use this function, submit a service ticket. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket**.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click its name. In the navigation pane on the left, choose **Backups & Restorations**.

Figure 8-18 Selecting a backup

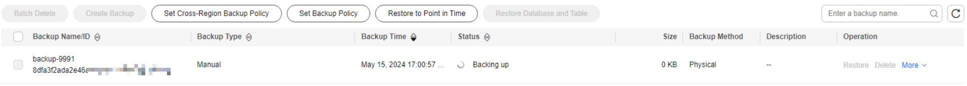

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
backup-9991 8dfa3f2ada2e46[redacted]	Manual	May 15, 2024 17:00:57 ...	Backing up	0 KB	Physical	--	Restore Delete More

Alternatively, on the navigation pane on the left, choose **Backups**.

Figure 8-19 Stopping a backup

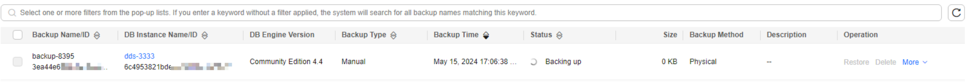
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all backup names matching this keyword.

Backup Name/ID	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
backup-8395 3ea44e6[redacted]	dds-3333 6c4953821bde[redacted]	Community Edition 4.4	Manual	May 15, 2024 17:06:38 ...	Backing up	0 KB	Physical	--	Restore Delete More

Step 5 Choose **More** > **Stop** in the **Operation** column.

Figure 8-20 Stopping a backup

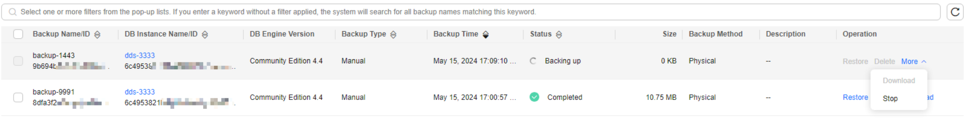
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all backup names matching this keyword.
Download
Stop

Backup Name/ID	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
backup-1443 9b694b[redacted]	dds-3333 6c49538[redacted]	Community Edition 4.4	Manual	May 15, 2024 17:09:10 ...	Backing up	0 KB	Physical	--	Restore Delete More
backup-9991 8dfa3f2[redacted]	dds-3333 6c4953821[redacted]	Community Edition 4.4	Manual	May 15, 2024 17:00:57 ...	Completed	10.75 MB	Physical	--	Restore

Step 6 On the displayed dialog box, click **Yes**.

----End

8.10 Downloading a Backup File

8.10.1 Using OBS Browser+

You can use OBS Brower+ to download a manual or an automated backup to a local device for backup or restoration.

Precautions

- When you use OBS Browser+ to download backup data, you will not be billed for outbound traffic from OBS.
- If the size of a backup file is greater than 400 MB, use OBS Browser+ to download the backup file.
- Backups downloaded from the DDS console are all full backups.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

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

Step 5 On the **Backups** page, click the **Clusters**, **Replica Sets**, or **Single Nodes** tab, locate the available backup you want to download and click **Download** in the **Operation** column.

Figure 8-21 Download Backup

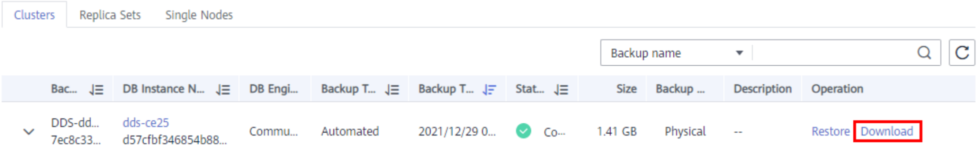

Bac...	DB Instance N...	DB Engi...	Backup T...	Backup T...	Stat...	Size	Backup ...	Description	Operation
DDS-dd... 7ec8c33...	dds-ce25 d57cfbf346854b88...	Commu...	Automated	2021/12/29 0...	✓ Co...	1.41 GB	Physical	--	Restore Download

Step 6 In the displayed dialog box, select **Use OBS Browser+** and click **OK**.

Figure 8-22 Selecting a download method

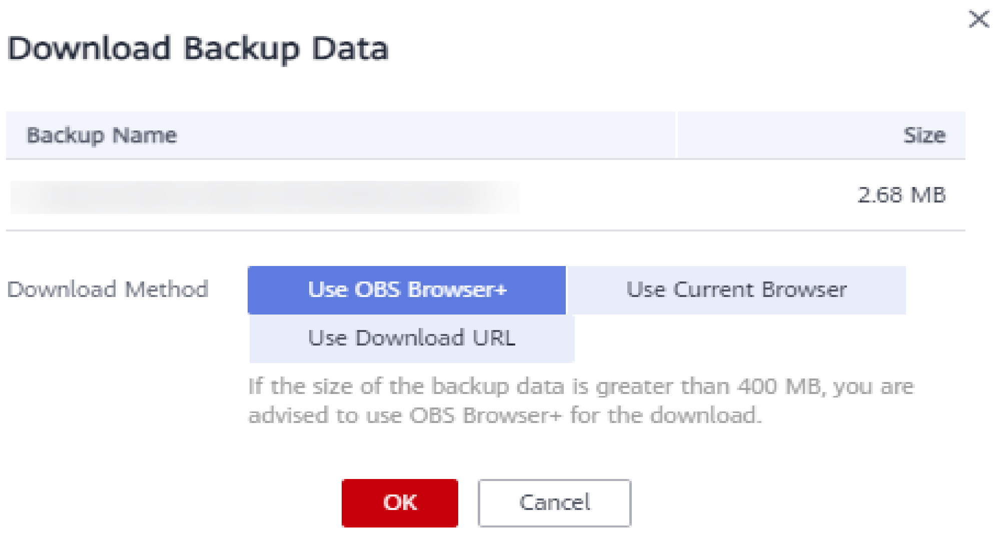

Download Backup Data

Backup Name	Size
[redacted]	2.68 MB

Download Method: Use OBS Browser+ | Use Current Browser | Use Download URL

If the size of the backup data is greater than 400 MB, you are advised to use OBS Browser+ for the download.

Step 7 On the displayed page, download the DDS backup file as prompted.

Figure 8-23 Download guide page

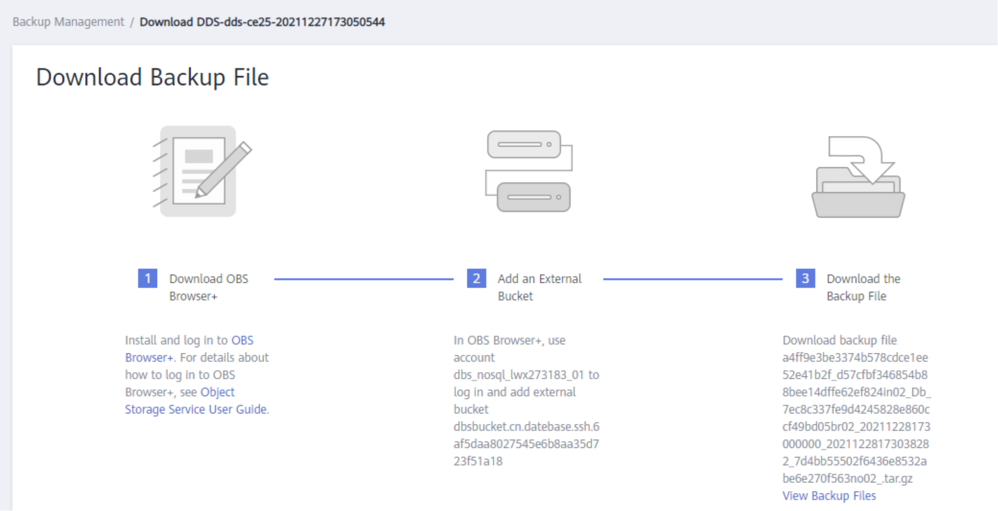

Backup Management / Download DDS-dds-ce25-20211227173050544

Download Backup File

1 Download OBS Browser+

Install and log in to OBS Browser+. For details about how to log in to OBS Browser+, see Object Storage Service User Guide.

2 Add an External Bucket

In OBS Browser+, use account dbs_nosql_lwx273183_01 to log in and add external bucket dbsbucket.cn.datebase.ssh.6af5daa8027545e6b8aa35d723f51a18

3 Download the Backup File

Download backup file a4ff9e3be3374b578cdce1ee52e41b2f_d57cfbf346854b88bee14dffe62ef824in02_Db_7ec8c337fe9d4245828e860ccf49bd05br02_20211228173000000_20211228173038282_7d4bb55502f6436e8532abe6e270f563no02_.tar.gz

View Backup Files

Step 8 Download OBS Browser+ following the step 1 provided in **Figure 8-23**.

Step 9 Decompress and install OBS Browser+.

Step 10 Log in to OBS Browser+.

For details about how to log in to OBS Browser+, see **Logging In to OBS Browser+** in the *Object Storage Service Tools Guide*.

Step 11 Add an external bucket.

In the **Add External Bucket** dialog box of OBS Browser+, enter the bucket name displayed in step 2 on page **Figure 8-23**, and click **OK**.

Step 12 Download the backup file.

On the OBS Browser+ page, click the external bucket that you added. In the search box on the right of OBS Browser+, enter the backup file name displayed in step 3 on page **Figure 8-23**. In the search result, locate the target backup and download it.

Step 13 After the backup file is downloaded, use the LZ4 to decompress the file.

Run the following command to decompress the backup file:

lz4 -d *$1* **| tar -xC** *$2*

$1: indicates the downloaded backup file.

$2: indicates the directory to which the backup file is decompressed.

Step 14 You can restore data locally as required.

For details, see the following documentation.

- **Restoring a Cluster Backup to an On-premises Database**
- **Restoring a Replica Set Backup to an On-Premises Database**

----End

8.10.2 Using Current Browser

You can user a browser to download a manual or an automated backup to a local device for backup or restoration

Precautions

- Cluster backup files cannot be downloaded using a browser.
- Backups downloaded from the DDS console are all full backups.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

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

Step 5 On the **Backups** page, click the **Clusters**, **Replica Sets**, or **Single Nodes** tab, locate the available backup you want to download and click **Download** in the **Operation** column.

Step 6 In the displayed dialog box, select **Use Current Browser** for **Download Method** and click **OK**.

Figure 8-24 Selecting a download method

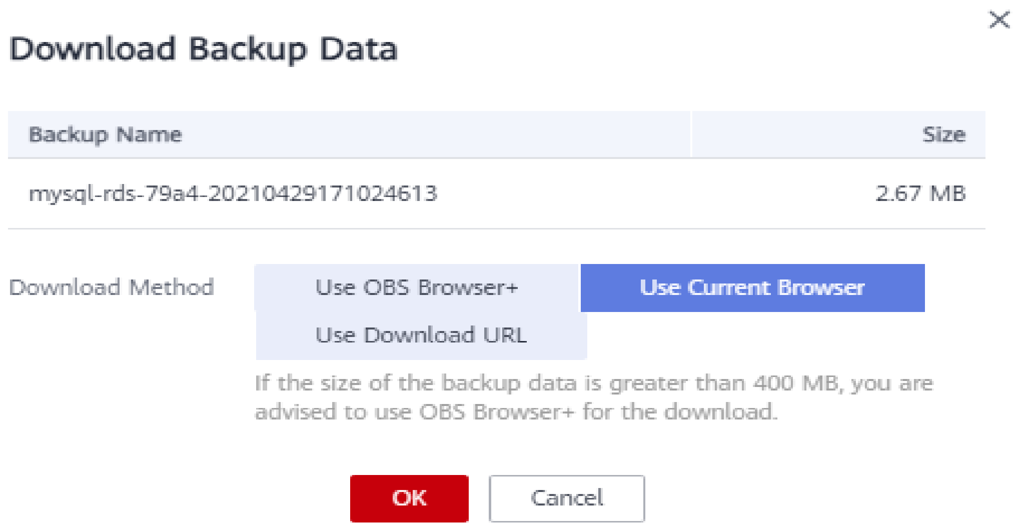

Download Backup Data

Backup Name	Size
mysql-rds-79a4-20210429171024613	2.67 MB

Download Method: Use Current Browser

If the size of the backup data is greater than 400 MB, you are advised to use OBS Browser+ for the download.

Step 7 After the backup file is downloaded, decompress it using LZ4.

Run the following command to decompress the backup file:

lz4 -d *$1* **| tar -xC** *$2*

$1: indicates the downloaded backup file.

$2: indicates the directory to which the backup file is decompressed.

Step 8 You can restore data locally as required.

For details, see the following documentation.

- **Restoring a Cluster Backup to an On-premises Database**
- **Restoring a Replica Set Backup to an On-Premises Database**

----End

8.10.3 Using Download URL

You can download manual or automated backup files using the URL provided by DDS to a local device for backup or restoration.

Precautions

Backups downloaded from the DDS console are all full backups.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

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

Step 5 On the **Backups** page, click the **Clusters**, **Replica Sets**, or **Single Nodes** tab, locate the available backup you want to download and click **Download** in the **Operation** column.

Step 6 In the displayed dialog box, select **Use Download URL** for **Download Method**, click ⧉ to copy the URL, and click **OK**.

Figure 8-25 Selecting a download method

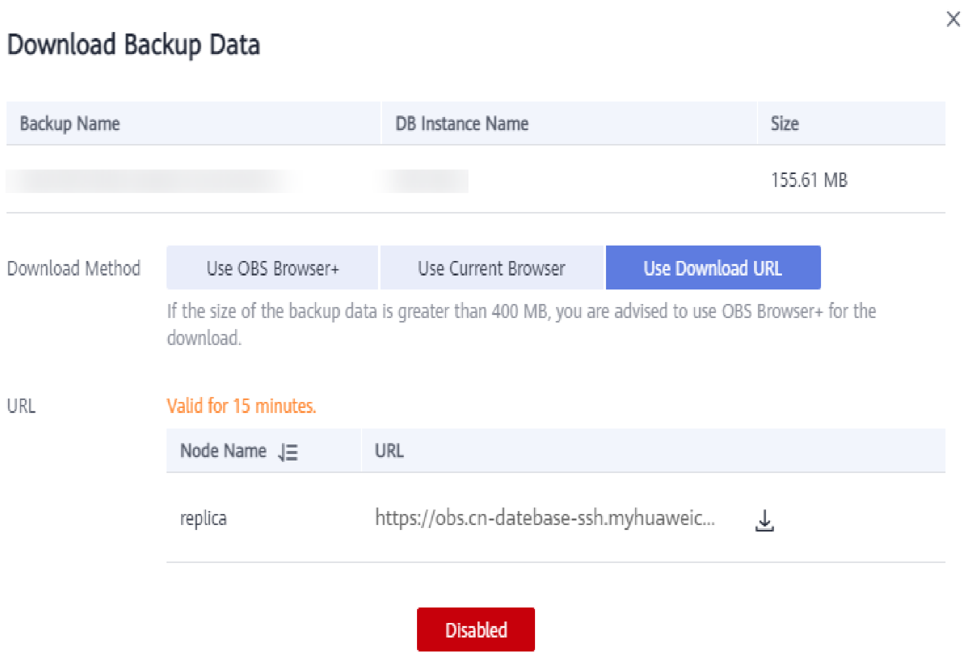

Download Backup Data

Backup Name	DB Instance Name	Size
[redacted]	[redacted]	155.61 MB

Download Method: Use OBS Browser+ | Use Current Browser | Use Download URL

If the size of the backup data is greater than 400 MB, you are advised to use OBS Browser+ for the download.

URL: Valid for 15 minutes.

Node Name	URL
replica	https://obs.cn-datebase-ssh.myhuaweic...

A valid URL for downloading the backup data is displayed.

- You can use various download tools, such as your browser and Xunlei to download backup files.
- You can also run the following command to download backup files:

  wget -O *FILE_NAME* **--no-check-certificate "***DOWNLOAD_URL***"**

  Parameter description:

  FILE_NAME is the new name of the downloaded backup file. The original backup file name may be too long and exceed the maximum characters allowed by the client file system, so you are advised to rename the backup file.

DOWNLOAD_URL is the location of the backup file to be downloaded. If the location contains special characters, escape is required.

Step 7 After the backup file is downloaded, decompress it using LZ4.

Run the following command to decompress the backup file:

lz4 -d *$1* **| tar -xC** *$2*

$1: indicates the downloaded backup file.

$2: indicates the directory to which the backup file is decompressed.

Step 8 You can restore data locally as required.

For details, see the following documentation.

- **Restoring a Cluster Backup to an On-premises Database**
- **Restoring a Replica Set Backup to an On-Premises Database**

----End

9 Data Restorations

9.1 Solutions

DDS provides multiple data restoration solutions. You can select a proper solution to meet your service requirements.

NOTE

By default, all DDS versions 3.2, 3.4, 4.0, 4.2, and 4.4 are supported unless otherwise specified.

Table 9-1 Solutions

Restoration Type	Instance Type and Version	Scenario
Restoring Data to a New Instance	- Cluster - Replica set - Single node	You can restore an existing automated or manual backup file to a new instance.
Restoring Data to the Original Instance	- Cluster - Replica set - Single node	You can restore an existing automated or manual backup file to the original instance.
Restoring Data to a Point in Time	- Cluster 4.0 or later - Replica set 4.0 or later	You can restore an instance to a point in time.
Restoring Database Tables to a Point in Time	- Replica set 4.0 or later - Cluster 4.0 or later	You can restore a database table to a point in time.

Restoration Type	Instance Type and Version	Scenario
Restoring Data to an On-Premises Database	● Cluster (versions 3.4 and 4.0) ● Replica set (versions 3.4 and 4.0) ● Single node (versions 3.4 and 4.0)	You can download a DDS backup file to your local PC and restore data to an on-premises database.
Restoring Data Using mongorestore	● Cluster ● Replica set ● Single node	You can use tools provided by the MongoDB client to restore data.
Restoring Data Using mongoimport	● Cluster ● Replica set ● Single node	You can use tools provided by the MongoDB client to restore data.

9.2 Restoring Data to a New Instance

9.2.1 Restoring a Cluster Backup to a New Instance

DDS allows you to restore an existing automated or manual backup to a new instance. The restored data is the same as the backup data.

When you restore an instance from a backup file, a full backup file is downloaded from OBS and then restored to the instance at an average speed of 40 MB/s.

Precautions

To restore backup files to a new instance, your account balance must be greater than or equal to $0 USD. You will pay for the new instance specifications.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name. Choose **Backups & Restorations** in the navigation pane, select the backup to be restored, and click **Restore**.

Figure 9-1 Restoring a cluster from a backup

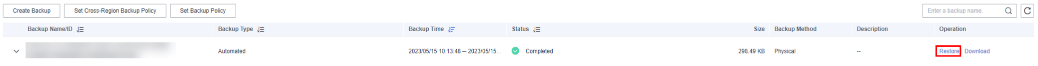

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
[redacted]	Automated	2023/05/15 10:13:48 – 2023/05/15...	Completed	298.49 KB	Physical	--	Restore Download

Alternatively, on the navigation pane on the left, choose **Backups**. On the **Backups** page, locate the target backup on the **Clusters** tab and click **Restore** in the **Operation** column.

Figure 9-2 Restoring a cluster from a backup

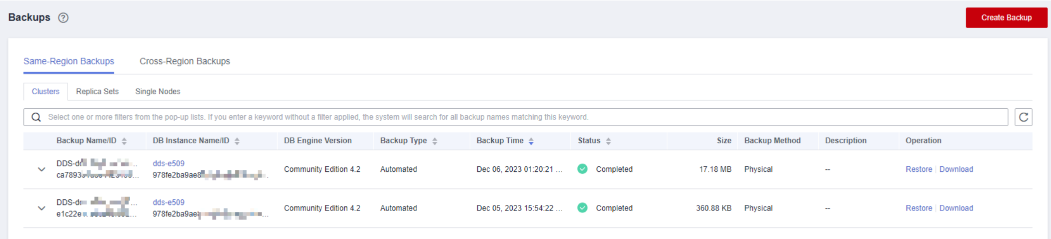
Backups
Create Backup
Same-Region Backups
Cross-Region Backups
Clusters
Replica Sets
Single Nodes
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all backup names matching this keyword.

Backup Name/ID	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-d[redacted] ca7893[redacted]	dds-e509 978fe2ba9ae8[redacted]	Community Edition 4.2	Automated	Dec 06, 2023 01:20:21 ...	Completed	17.18 MB	Physical	--	Restore | Download
DDS-d[redacted] e1c22e[redacted]	dds-e509 978fe2ba9ae[redacted]	Community Edition 4.2	Automated	Dec 05, 2023 15:54:22 ...	Completed	360.88 KB	Physical	--	Restore | Download

Step 5 In the **Restore DB Instance** dialog box, select **Create New Instance** for **Restoration Method** and click **OK**.

Figure 9-3 Restoring a cluster backup to a new instance

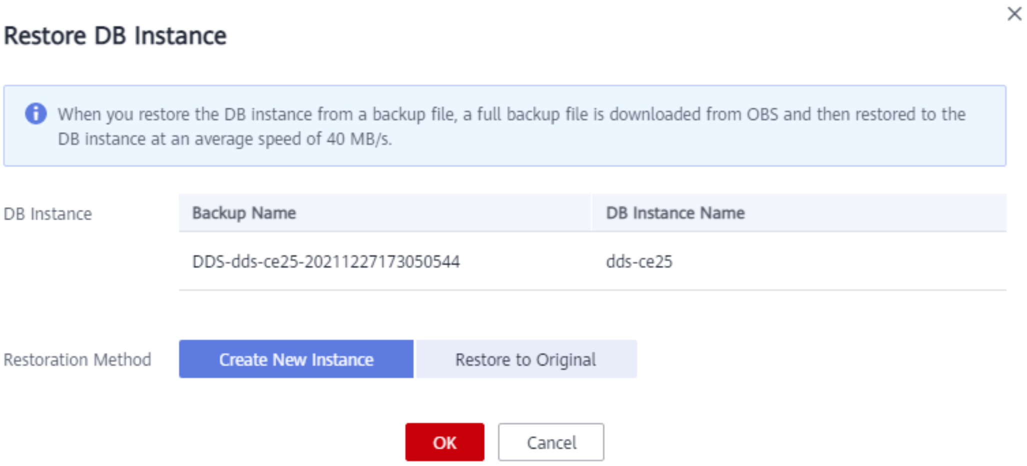

When you restore the DB instance from a backup file, a full backup file is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

DB Instance

Backup Name	DB Instance Name
DDS-dds-ce25-20211227173050544	dds-ce25

Restoration Method: Create New Instance

Step 6 The **Create New Instance** page is displayed for you to create an instance using the backup data. The new DB instance is independent from the original one.

- You are recommended to deploy the restored DB instance in a different AZ to ensure that applications will not be adversely affected by the failure in any single AZ.
- The database type, DB instance type, compatible MongoDB version, storage engine, storage type, and shard quantity must be the same as those of the original and cannot be changed.
- The number of dds mongos nodes is 2 by default and ranges from 2 to 16. You can specify the quantity.
- The storage space is the same as that of the original shard node by default. You can increase the storage space, but you cannot reduce it.
- Other settings are the same as those of the original DB instance by default and can be modified. For details, see **Buying a Cluster Instance**.

- A full backup is triggered after the new instance is created.

----End

9.2.2 Restoring a Replica Set Backup to a New Instance

DDS allows you to restore an existing automated or manual backup to a new instance. The restored data is the same as the backup data.

When you restore an instance from a backup file, a full backup file is downloaded from OBS and then restored to the instance at an average speed of 40 MB/s.

Precautions

To restore backup files to a new instance, your account balance must be greater than or equal to $0 USD. You will pay for the new instance specifications.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance. Choose **Backups & Restorations** in the navigation pane, select the backup to be restored, and click **Restore**.

Figure 9-4 Restoring a replica set instance backup

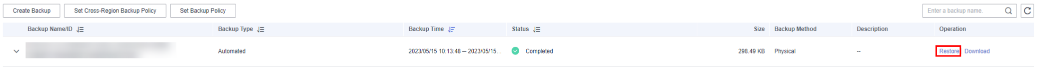

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
[redacted]	Automated	2023/05/15 10:13:48 – 2023/05/15...	Completed	298.49 KB	Physical	--	Restore Download

Alternatively, on the navigation pane on the left, choose **Backups**. On the **Backups** page, locate the backup on the **Replica Sets** tab and click **Restore** in the **Operation** column.

Figure 9-5 Restoring a replica set instance backup

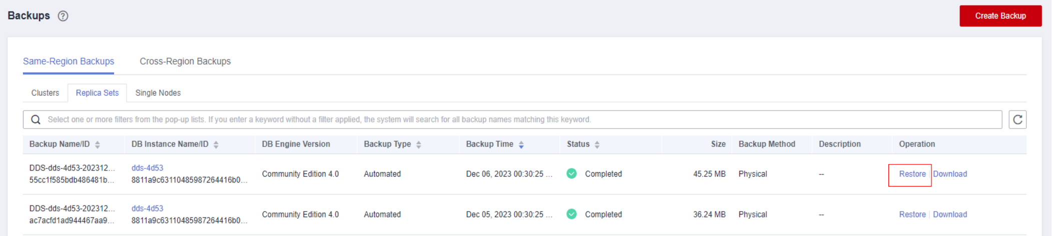
Backups
Create Backup
Same-Region Backups
Cross-Region Backups
Clusters
Replica Sets
Single Nodes
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all backup names matching this keyword.

Backup Name/ID	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dds-4d53-202312... 55cc1f585bdb486481b...	dds-4d53 8811a9c63110485987264416b0...	Community Edition 4.0	Automated	Dec 06, 2023 00:30:25 ...	Completed	45.25 MB	Physical	--	Restore Download
DDS-dds-4d53-202312... ac7acfd1ad944467aa9...	dds-4d53 8811a9c63110485987264416b0...	Community Edition 4.0	Automated	Dec 05, 2023 00:30:25 ...	Completed	36.24 MB	Physical	--	Restore Download

Step 5 In the **Restore DB Instance** dialog box, select **Create New Instance** for **Restoration Method** and click **OK**.

Figure 9-6 Restoring to a new instance

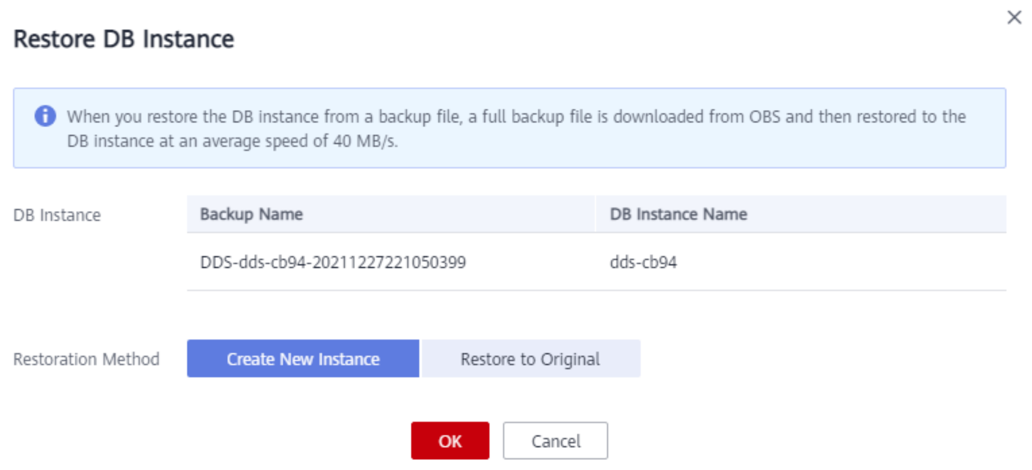

Restore DB Instance

When you restore the DB instance from a backup file, a full backup file is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

DB Instance

Backup Name	DB Instance Name
DDS-dds-cb94-20211227221050399	dds-cb94

Restoration Method: Create New Instance | Restore to Original

Step 6 The **Create New Instance** page is displayed for you to create an instance using the backup data. The new DB instance is independent from the original one.

- You are recommended to deploy the restored DB instance in a different AZ to ensure that applications will not be adversely affected by the failure in any single AZ.
- The database type, DB instance type, compatible MongoDB version, storage engine, and storage type must be the same as those of the original and cannot be changed.
- The storage space is the same as that of the original instance by default. You can increase the storage space, but you cannot reduce it.
- Other settings have default values and can be modified. For details, see **Buying a Replica Set Instance**.
- A full backup is triggered after the new instance is created.

----End

9.2.3 Restoring a Single Node Backup to a New Instance

DDS allows you to restore an existing automated or manual backup to a new instance. The restored data is the same as the backup data.

When you restore an instance from a backup file, a full backup file is downloaded from OBS and then restored to the instance at an average speed of 40 MB/s.

NOTE

Huawei Cloud has discontinued the sale of DDS single node instances since July 15, 2023.

Precautions

To restore backup files to a new instance, your account balance must be greater than or equal to $0 USD. You will pay for the new instance specifications.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the single node instance name. Choose **Backups & Restorations** in the navigation pane, select the backup to be restored, and click **Restore**.

Figure 9-7 Restoring a single node backup

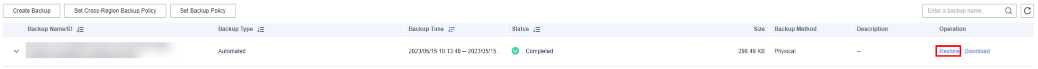

	Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
∨	[redacted]	Automated	2023/05/15 10:13:48 – 2023/05/15...	✓ Completed	298.49 KB	Physical	--	Restore Download

Alternatively, on the navigation pane on the left, choose **Backups**. On the **Backups** page, locate the target backup on the **Single Nodes** tab and click **Restore** in the **Operation** column.

Figure 9-8 Restoring a single node backup

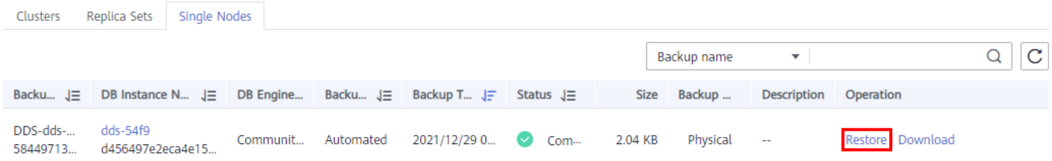

Backu...	DB Instance N...	DB Engine...	Backu...	Backup T...	Status	Size	Backup ...	Description	Operation
DDS-dds-... 58449713...	dds-54f9 d456497e2eca4e15...	Communit...	Automated	2021/12/29 0...	✓ Com...	2.04 KB	Physical	--	Restore Download

Step 5 In the **Restore DB Instance** dialog box, select **Create New Instance** for **Restoration Method** and click **OK**.

Figure 9-9 Restoring a single node backup to a new instance

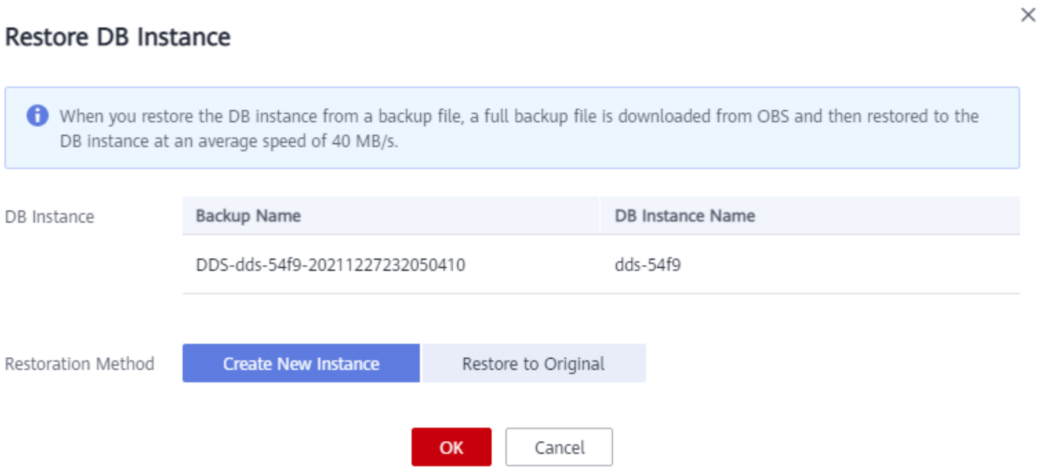
Restore DB Instance

When you restore the DB instance from a backup file, a full backup file is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

DB Instance

Backup Name	DB Instance Name
DDS-dds-54f9-20211227232050410	dds-54f9

Restoration Method: Create New Instance | Restore to Original

Step 6 The **Create New Instance** page is displayed for you to create an instance using the backup data. The new DB instance is independent from the original one.

- You are recommended to deploy the restored DB instance in a different AZ to ensure that applications will not be adversely affected by the failure in any single AZ.

- The database type, DB instance type, compatible MongoDB version, storage engine, and storage type must be the same as those of the original and cannot be changed.
- The storage space is the same as that of the original instance by default. You can increase the storage space, but you cannot reduce it.
- A full backup is triggered after the new instance is created.

----End

9.2.4 Restoring a Cross-Region Backup to a New DB Instance

DDS allows you to restore an existing automated backup to a new instance. The restored data is the same as the backup data.

When you restore an instance from a backup file, a full backup file is downloaded from OBS and then restored to the instance at an average speed of 40 MB/s.

Precautions

To restore backup files to a new instance, your account balance must be greater than or equal to $0 USD. You will pay for the new instance specifications.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, choose **Backups** on the navigation pane. Click the **Cross-Region Backups** tab. On the displayed page, locate the target DB instance and click **View Cross-Region Backup** in the **Operation** column.

Figure 9-10 Cross-region backups

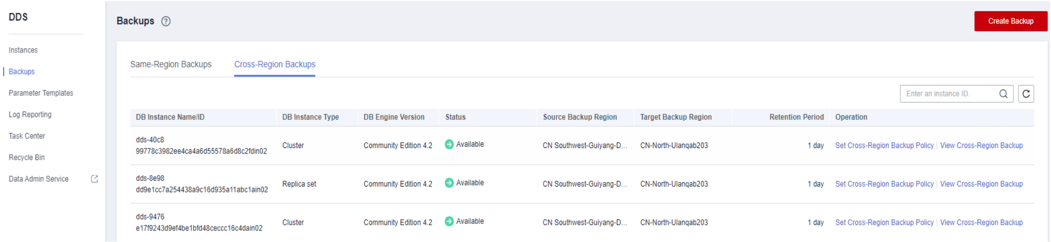

DB Instance Name/ID	DB Instance Type	DB Engine Version	Status	Source Backup Region	Target Backup Region	Retention Period	Operation
dds-40c8 99778c3982ee4ca4a6d55578a6d8c2fdin02	Cluster	Community Edition 4.2	Available	CN Southwest-Guiyang-D...	CN-North-Ulanqab203	1 day	Set Cross-Region Backup Policy | View Cross-Region Backup
dds-8e98 dd9e1cc7a254438a9c16d935a11abc1ain02	Replica set	Community Edition 4.2	Available	CN Southwest-Guiyang-D...	CN-North-Ulanqab203	1 day	Set Cross-Region Backup Policy | View Cross-Region Backup
dds-9476 e17f9243d9ef4be1bfd48ceccc16c4dain02	Cluster	Community Edition 4.2	Available	CN Southwest-Guiyang-D...	CN-North-Ulanqab203	1 day	Set Cross-Region Backup Policy | View Cross-Region Backup

Step 5 Locate the backup to be restored and click **Restore** in the **Operation** column.

Figure 9-11 Restoring a cross-region backup

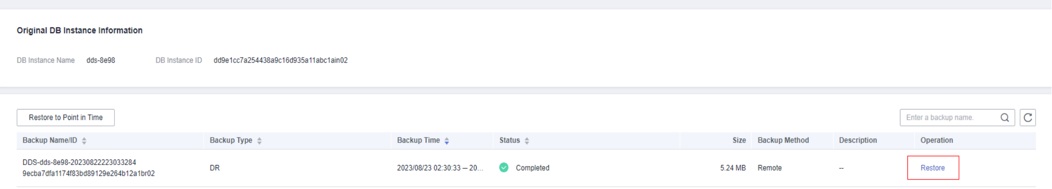

Original DB Instance Information

DB Instance Name dds-8e98 DB Instance ID dd9e1cc7a254438a9c16d935a11abc1ain02

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dds-8e98-20230822223033284 9ecba7dfa1174f83bd89129e264b12a1br02	DR	2023/08/23 02:30:33 – 20...	Completed	5.24 MB	Remote	--	Restore

Step 6 In the **Restore DB Instance** dialog box, select **Create New Instance** for **Restoration Method** and click **OK**.

Figure 9-12 Restoring a cross-region backup to a new DB instance

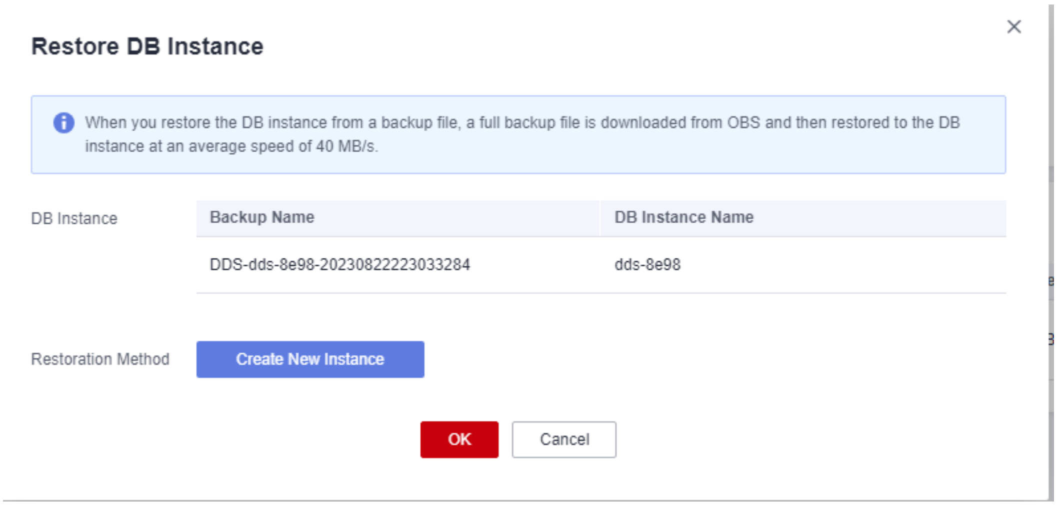

Restore DB Instance

When you restore the DB instance from a backup file, a full backup file is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

DB Instance

Backup Name	DB Instance Name
DDS-dds-8e98-20230822223033284	dds-8e98

Restoration Method: Create New Instance

Step 7 The **Create New Instance** page is displayed for you to create an instance using the backup data. The new DB instance is independent from the original one.

- You are recommended to deploy the restored DB instance in a different AZ to ensure that applications will not be adversely affected by the failure in any single AZ.
- The database type, DB instance type, compatible MongoDB version, storage engine, storage type, and shard quantity must be the same as those of the original and cannot be changed.
- The storage space is the same as that of the original instance by default. You can increase the storage space, but you cannot reduce it.
- Other settings are the same as those of the original DB instance by default and can be modified. For details, see **Buying a Cluster Instance** or **Buying a Replica Set Instance**.
- A full backup is triggered after the new instance is created.

----End

9.3 Restoring Data to the Original Instance

9.3.1 Restoring a Cluster Backup to the Original Instance

DDS allows you to restore an existing automated or manual backup to an original instance. The restored data is the same as the backup data.

When you restore an instance from a backup file, a full backup file is downloaded from OBS and then restored to the instance at an average speed of 40 MB/s.

Precautions

- Restoring backup data to the original instance will overwrite existing data on the instance and cause the instance to be unavailable during the restoration. Exercise caution when performing this operation.
- The administrator password of the instance remains unchanged after the restoration.
- If you restore a manual backup, check whether the instance to which the manual backup belongs exists. If the instance does not exist, the backup can only be restored to a new instance.
- If a cluster DB instance have read replicas associated, backup data can only be restored to a new DB instance.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name. Choose **Backups & Restorations** in the navigation pane on the left, select the backup to be restored, and click **Restore**.

Figure 9-13 Restoring a cluster from a backup

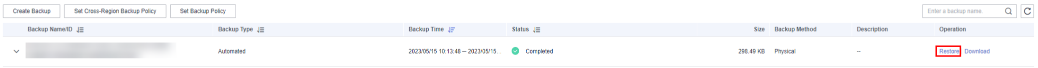

	Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
⌄	[redacted]	Automated	2023/05/15 10:13:48 – 2023/05/15...	✓ Completed	298.49 KB	Physical	--	Restore Download

Alternatively, on the navigation pane on the left, choose **Backups**. On the **Backups** page, locate the target backup on the **Clusters** tab and click **Restore** in the **Operation** column.

Figure 9-14 Restoring a cluster from a backup

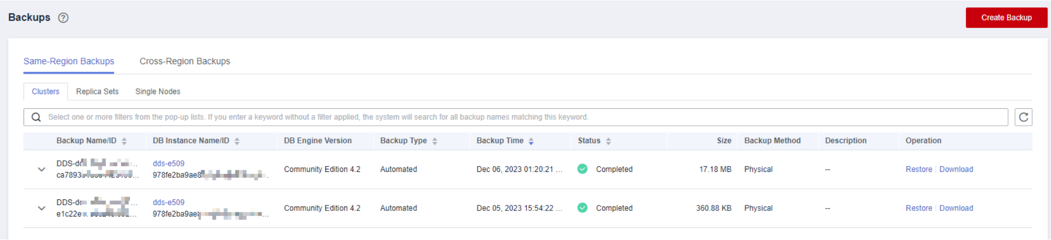
Backups
Create Backup
Same-Region Backups
Cross-Region Backups
Clusters
Replica Sets
Single Nodes
Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all backup names matching this keyword.

	Backup Name/ID	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
⌄	DDS-d[redacted] ca7893[redacted]	dds-e509 978fe2ba9ae8[redacted]	Community Edition 4.2	Automated	Dec 06, 2023 01:20:21 ...	✓ Completed	17.18 MB	Physical	--	Restore | Download
⌄	DDS-d[redacted] e1c22e[redacted]	dds-e509 978fe2ba9ae[redacted]	Community Edition 4.2	Automated	Dec 05, 2023 15:54:22 ...	✓ Completed	360.88 KB	Physical	--	Restore | Download

Step 5 In the **Restore DB Instance** dialog box, select **Restore to Original** for **Restoration Method** and click **OK**.

Figure 9-15 Restore to Original

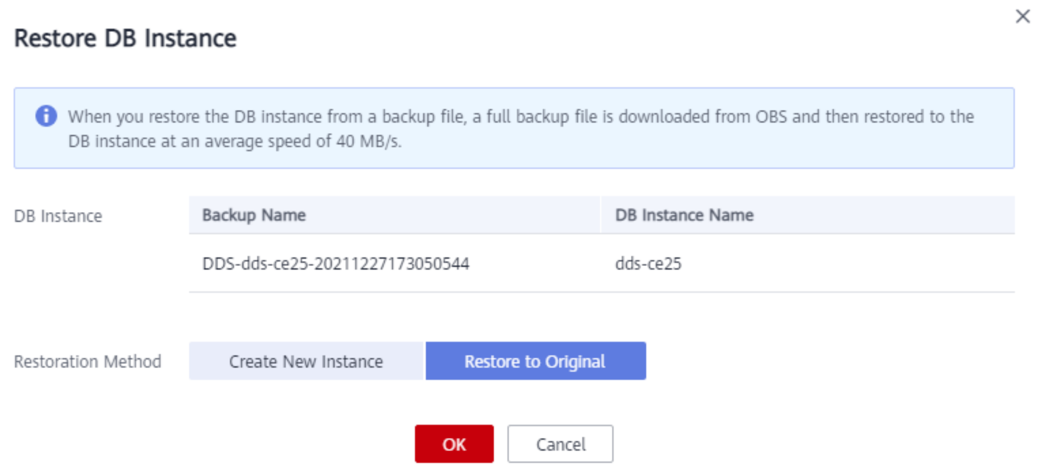

Restore DB Instance

When you restore the DB instance from a backup file, a full backup file is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

DB Instance

Backup Name	DB Instance Name
DDS-dds-ce25-20211227173050544	dds-ce25

Restoration Method: Create New Instance | Restore to Original

- On the **Instances** page, the status of the instance changes from **Restoring** to **Available**.
- After the restoration is complete, a full backup will be automatically triggered.

----End

9.3.2 Restoring a Replica Set Backup to the Original Instance

DDS allows you to restore an existing automated or manual backup to an original instance. The restored data is the same as the backup data.

When you restore an instance from a backup file, a full backup file is downloaded from OBS and then restored to the instance at an average speed of 40 MB/s.

Precautions

- Restoring backup data to the original instance will overwrite existing data on the instance and cause the instance to be unavailable during the restoration. Exercise caution when performing this operation.
- The administrator password of the instance remains unchanged after the restoration.
- If you restore a manual backup, check whether the instance to which the manual backup belongs exists. If the instance does not exist, the backup can only be restored to a new instance.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance. Choose **Backups & Restorations** in the navigation pane on the left, select the backup to be restored, and click **Restore**.

Figure 9-16 Restoring a replica set instance backup

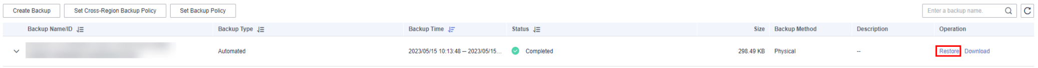

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
[redacted]	Automated	2023/05/15 10:13:48 – 2023/05/15...	Completed	298.49 KB	Physical	--	Restore Download

Alternatively, on the navigation pane on the left, choose **Backups**. On the **Backups** page, locate the backup on the **Replica Sets** tab and click **Restore** in the **Operation** column.

Figure 9-17 Restoring a replica set instance backup

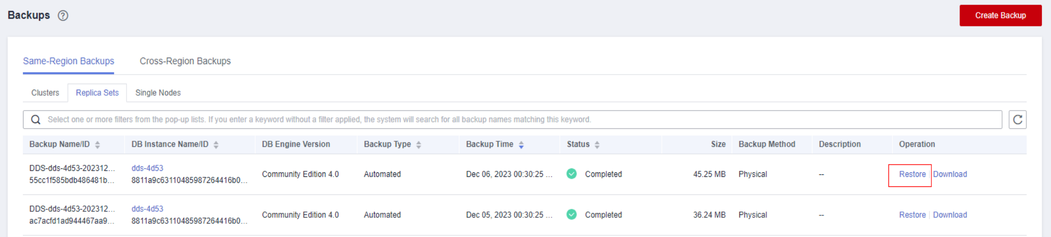

Select one or more filters from the pop-up lists. If you enter a keyword without a filter applied, the system will search for all backup names matching this keyword.

Backup Name/ID	DB Instance Name/ID	DB Engine Version	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dds-4d53-202312... 55cc1f585bdb486481b...	dds-4d53 8811a9c63110485987264416b0...	Community Edition 4.0	Automated	Dec 06, 2023 00:30:25 ...	Completed	45.25 MB	Physical	--	Restore Download
DDS-dds-4d53-202312... ac7acfd1ad944467aa9...	dds-4d53 8811a9c63110485987264416b0...	Community Edition 4.0	Automated	Dec 05, 2023 00:30:25 ...	Completed	36.24 MB	Physical	--	Restore Download

Step 5 In the **Restore DB Instance** dialog box, select **Restore to Original** for **Restoration Method** and click **OK**.

Figure 9-18 Restore to Original

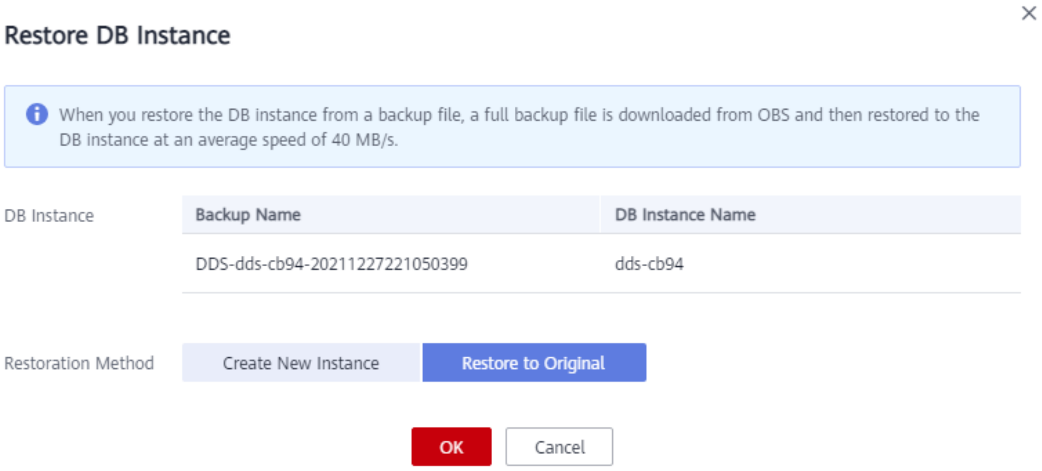

When you restore the DB instance from a backup file, a full backup file is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

DB Instance

Backup Name	DB Instance Name
DDS-dds-cb94-20211227221050399	dds-cb94

- On the **Instances** page, the status of the instance changes from **Restoring** to **Available**.
- After the restoration is complete, a full backup will be automatically triggered.

----End

9.3.3 Restoring a Single Node Backup to the Original Instance

DDS allows you to restore an existing automated or manual backup to an original instance. The restored data is the same as the backup data.

When you restore an instance from a backup file, a full backup file is downloaded from OBS and then restored to the instance at an average speed of 40 MB/s.

NOTE

Huawei Cloud has discontinued the sale of DDS single node instances since July 15, 2023.

Precautions

- Restoring backup data to the original instance will overwrite existing data on the instance and cause the instance to be unavailable during the restoration. Exercise caution when performing this operation.
- The administrator password of the instance remains unchanged after the restoration.
- If you restore a manual backup, check whether the instance to which the manual backup belongs exists. If the instance does not exist, the backup can only be restored to a new instance.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the single node instance name. Choose **Backups & Restorations** in the navigation pane on the left, select the backup to be restored, and click **Restore**.

Figure 9-19 Restoring a single node backup

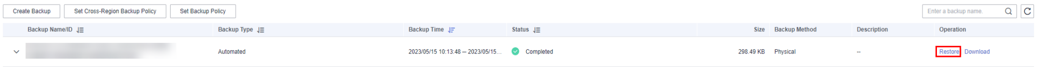

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
[redacted]	Automated	2023/05/15 10:13:48 — 2023/05/15...	Completed	298.49 KB	Physical	--	Restore Download

Alternatively, on the navigation pane on the left, choose **Backups**. On the **Backups** page, locate the target backup on the **Single Nodes** tab and click **Restore** in the **Operation** column.

Figure 9-20 Restoring a single node backup

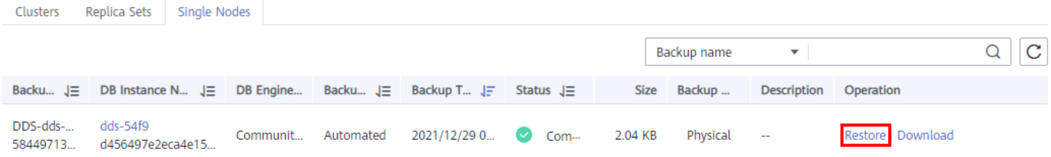

Backu...	DB Instance N...	DB Engine...	Backu...	Backup T...	Status	Size	Backup ...	Description	Operation
DDS-dds-... 58449713...	dds-54f9 d456497e2eca4e15...	Communit...	Automated	2021/12/29 0...	Com...	2.04 KB	Physical	--	Restore Download

Step 5 In the **Restore DB Instance** dialog box, select **Restore to Original** for **Restoration Method** and click **OK**.

Figure 9-21 Restore to Original

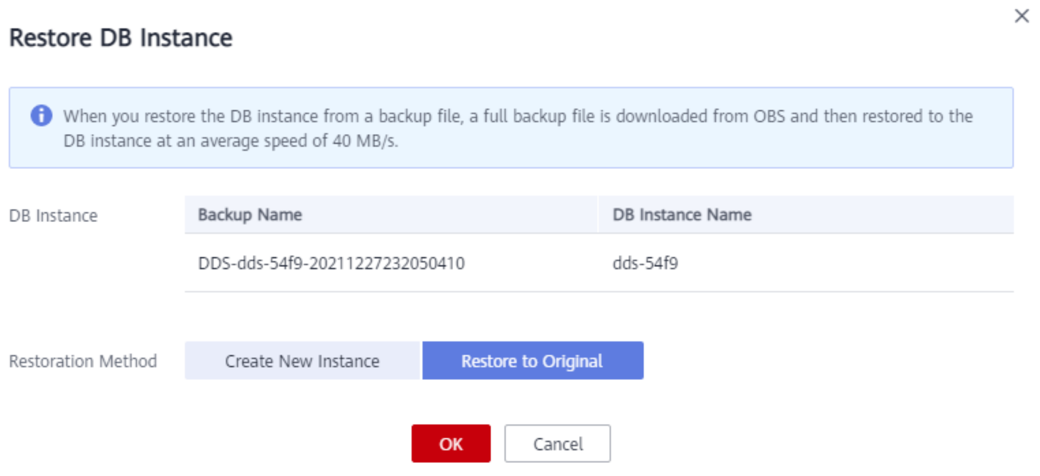

Restore DB Instance

When you restore the DB instance from a backup file, a full backup file is downloaded from OBS and then restored to the DB instance at an average speed of 40 MB/s.

DB Instance

Backup Name	DB Instance Name
DDS-dds-54f9-20211227232050410	dds-54f9

Restoration Method: Create New Instance / Restore to Original

- On the **Instances** page, the status of the instance changes from **Restoring** to **Available**.
- After the restoration is complete, a full backup will be automatically triggered.

----End

9.4 Restoring Data to a Point in Time

9.4.1 Restoring a Cluster Instance to a Point in Time

DDS allows you to restore cluster instances to a point in time.

When you enter the point in time that you want to restore the instance to, DDS downloads the most recent full backup file from OBS to the instance. Then, incremental backups are also restored to the specified point in time on the instance. Data is restored at an average speed of 30 MB/s.

Precautions

- To use this function, contact customer service to apply for the corresponding permission.
- Only cluster instances of version 4.0 or later can be restored to a specified point in time.
- Data can be restored to a specific point in time only after the automated and incremental backup policies are enabled.
- Data can be restored to a new instance or the original instance.
- To ensure data security, the dropDatabase operation is blocked when the incremental backup is restored to a point in time. Empty databases or views may exist after the restoration. You can delete them.
- Data cannot be restored to a point in time in any of the following scenarios: rename operation, collmod operation, creating a user, deleting a user, creating a role, deleting a role, enabling shard IP addresses of a cluster instance,

changing the password of the shard node user, enabling config IP addresses of a cluster instance, changing the password of the config node user, and changing the password of the **rwuser** user. When a restricted scenario occurs, the incremental backup stops. After the next automated full backup, the incremental backup resumes.

- If the time window of the full backup overlaps with that of the incremental backup, the full backup prefers. The incremental backup is restricted so that a few time ranges are not within the recovery time window.
- If the incremental oplog traffic is greater than 250 GB/h, the incremental backup speed may not keep up with the oplog generation speed. As a result, some restoration time points may be unavailable.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

If you want compute and network resources dedicated to your exclusive use, **enable a DeC** and **apply for DCC resources**. After enabling a DeC, you can select the DeC region and project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Restore to Point in Time**.

Figure 9-22 Restoring a cluster instance to a point in time

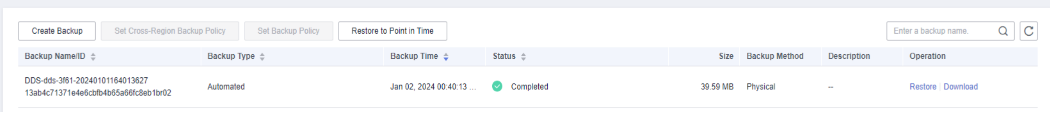

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dds-3f61-20240101164013627 13ab4c71371e4e6cbfb4b65a66fc8eb1br02	Automated	Jan 02, 2024 00:40:13 ...	Completed	39.59 MB	Physical	--	Restore | Download

Step 7 Select the date and time range, select or enter a time point within the acceptable range, and select **Create New Instance** or **Restore to Original**.

Figure 9-23 Restoring to a point in time

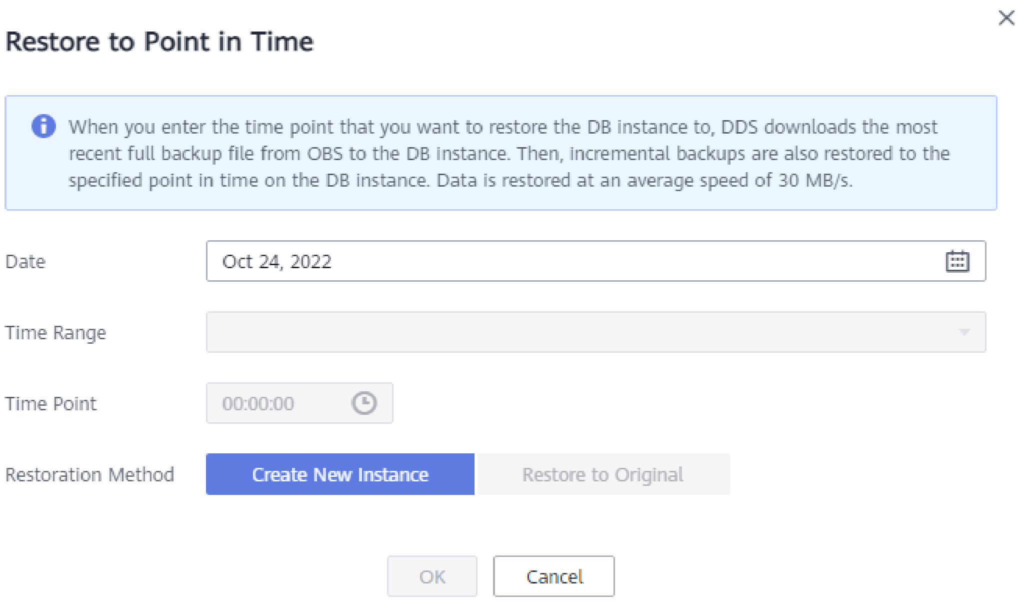

Restore to Point in Time

When you enter the time point that you want to restore the DB instance to, DDS downloads the most recent full backup file from OBS to the DB instance. Then, incremental backups are also restored to the specified point in time on the DB instance. Data is restored at an average speed of 30 MB/s.

Date: Oct 24, 2022

Time Range:

Time Point: 00:00:00

Restoration Method: Create New Instance | Restore to Original

Step 8 On the displayed page, the instance is restored based on the restoration method you selected in **Step 7**.

- Create New Instance

  The **Create New Instance** page is displayed for you to create an instance using the backup data. The new instance is independent from the original one.

  – You are recommended to deploy the restored instance in a different AZ to ensure that applications will not be adversely affected by the failure in any single AZ.

  – The database type, DB instance type, compatible MongoDB version, storage engine, and storage type must be the same as those of the original and cannot be changed.

  – The storage space is the same as that of the original instance by default. You can increase the storage space, but you cannot reduce it.

  – Other settings can be modified. For details, see **Buying a Cluster Instance**.

- Restore to Original

  Check that the status of the instance on the **Instances** page is **Restoring**.

NOTICE

– Restoring backup data to the original instance will overwrite existing data on the instance and cause the instance to be unavailable during the restoration. Exercise caution when performing this operation.

– The administrator password of the instance remains unchanged after the restoration.

----End

9.4.2 Restoring a Replica Set Instance to a Point in Time

You can restore a replica set instance to a specific point in time.

When you enter the point in time that you want to restore the instance to, DDS downloads the most recent full backup file from OBS to the instance. Then, incremental backups are also restored to the specified point in time on the instance. Data is restored at an average speed of 30 MB/s.

Precautions

- Currently, you can restore a replica set instance to a new or original DB instance at a point in time.
- Only replica set instances of version 4.0 or later can be restored to a point in time.
- Data can be restored to a specific point in time only after the automated backup policy is enabled.
- The local database is not included in the databases that can be restored to a specified time point.
- To ensure data security, the dropDatabase operation is blocked when the incremental backup is restored to a point in time. Empty databases or views may exist after the restoration. You can delete them.
- If the incremental oplog traffic is greater than 250 GB/h, the incremental backup speed may not keep up with the oplog generation speed. As a result, some restoration time points may be unavailable.
- Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Procedure

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

Step 2 Click ⦿ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance name.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Restore to Point in Time**.

Figure 9-24 Restoring to a point in time

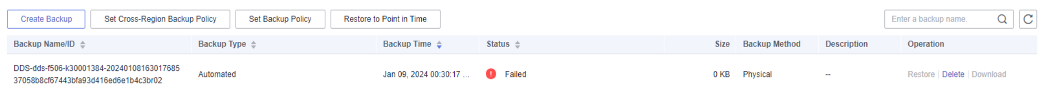

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dds-f506-k30001384-20240108163017685 37058b8cf67443bfa93d416ed6e1b4c3br02	Automated	Jan 09, 2024 00:30:17 ...	● Failed	0 KB	Physical	--	Restore | Delete | Download

Step 7 Select the date and time range, select or enter a time point within the acceptable range, and select **Create New Instance** or **Restore to Original**.

Figure 9-25 Restoring to a point in time

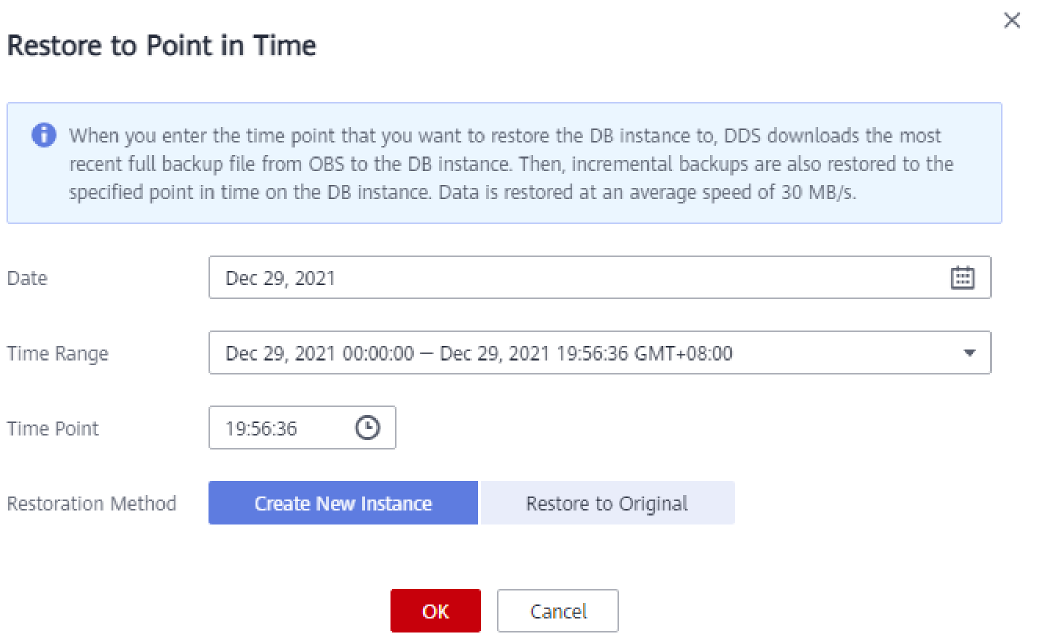

Restore to Point in Time

When you enter the time point that you want to restore the DB instance to, DDS downloads the most recent full backup file from OBS to the DB instance. Then, incremental backups are also restored to the specified point in time on the DB instance. Data is restored at an average speed of 30 MB/s.

Date: Dec 29, 2021

Time Range: Dec 29, 2021 00:00:00 – Dec 29, 2021 19:56:36 GMT+08:00

Time Point: 19:56:36

Restoration Method: Create New Instance | Restore to Original

Step 8 On the displayed page, the DB instance is restored based on the restoration method you selected in **Step 7**.

- Create New Instance

  The **Create New Instance** page is displayed for you to create an instance using the backup data. The new DB instance is independent from the original one.

  – You are recommended to deploy the restored DB instance in a different AZ to ensure that applications will not be adversely affected by the failure in any single AZ.

  – The database type, DB instance type, compatible MongoDB version, storage engine, and storage type must be the same as those of the original and cannot be changed.

  – The storage space is the same as that of the original instance by default. You can increase the storage space, but you cannot reduce it.

  – Other settings have default values and can be modified. For details, see **Buying a Replica Set Instance**.

- Restore to Original

NOTICE

– Restoring backup data to the original instance will overwrite existing data on the instance and cause the instance to be unavailable during the restoration. Exercise caution when performing this operation.

– The administrator password of the instance remains unchanged after the restoration.

– If the backup method is logical backup, the backup cannot be restored to the original instance.

Check that the status of the DB instance on the **Instances** page is **Restoring**.

----End

9.4.3 Restoring a Replica Set Database and Table to a Point in Time

To ensure data integrity and reduce impact on the original instance performance, the system restores the full and incremental data at the selected time point to a temporary instance, automatically exports the databases and tables to be restored, and then restores the databases and tables to the original instance. The time required depends on the amount of data to be backed up and restored on the instance. Please wait.

Restoring databases and tables will not overwrite data in the instance. You can select databases and tables to be restored.

Precautions

- Currently, only replica set instances of version 4.0 or later support the point-in-time recovery at the database and table level.
- Before performing the restoration, you need enable the automated backup policy.
- After a successful restoration, a new table named ***Original table name*****_bak_*****Timestamp*** is generated in the instance by default. If the table contains an index, the namespace of the index is changed to ***Original database name.Original table name*****_bak_*****Timestamp***. You can rename the table later as required.
- New databases and tables will be generated in the original DB instance. Ensure that sufficient storage space is available.
- The length of *<Database name>.<Table name>* cannot exceed 120 characters. The length of *<Database name>.<Table name>.<Index name>* cannot exceed 128 characters, or the restoration may fail.
- Ensure that the name of the restored table is different from that of the existing table, or the restoration may fail.
- If you perform a table-level restoration and the table does not exist at the required point in time, an empty table is automatically created. If you perform a database-level restoration, the missing table is not created.
- If the incremental oplog traffic is greater than 250 GB/h, the incremental backup speed may not keep up with the oplog generation speed. As a result, some restoration time points may be unavailable.
- Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Restrictions

- Database- and table-level restoration is related to CPU and memory specifications. For details about the maximum size of a single data record that can be restored to a specified time point, see **Table 9-2**.

- If a database- and table-level restoration fails, you can upgrade the specifications or batch restore the databases and tables to a specified time point.

Table 9-2 Specifications

CPU Type	Specifications	vCPUs	Memory (GB)	Maximum Size of a Single Data Record That Can Be Restored in a Collection
x86	General-purpose	2	4	400 KB
		2	8	800 KB
		4	8	1 MB
		4	16	1.3 MB
		8	16	1.3 MB
		8	32	2 MB
	Enhanced II	1	8	400 KB
		2	8	800 KB
		2	16	800 KB
		4	16	1.3 MB
		4	32	1.3 MB
		8	32	2 MB
		8	64	3 MB
		16	64	4 MB
		16	128	7 MB
		32	128	7 MB
		32	256	10 MB
		64	256	10 MB
		64	512	16 MB
Kunpeng	-	2	4	400 KB
	-	2	8	800 KB
	-	4	8	1 MB
	-	4	16	1.3 MB
	-	8	16	1.3 MB

CPU Type	Specifications	vCPUs	Memory (GB)	Maximum Size of a Single Data Record That Can Be Restored in a Collection
	-	8	32	2 MB
	-	16	32	2 MB
	-	16	64	4 MB

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the replica set instance.

Step 5 In the navigation pane on the left, choose **Backups & Restorations**.

Step 6 On the **Backups & Restorations** page, click **Restore Database and Table**.

Step 7 In the displayed dialog box, configure parameters as required. The data in the new database and table is the same as that in the database and table at the selected time point.

Table 9-3 Database information

Parameter	Description
Date	Date when the automated backup of the DB instance is generated.
Time Range	Time range during which the automated backup can be restored.
Time Point	The specific point in time when the automated full backup is generated.
Base Time Range	Time range during which the database and table can be restored based on the automated full backup.
Database and Table	Databases and tables that have been automatically backed up within the base time range are displayed on the left. Select the databases and tables on the left to sync information to the area on the right.
Time Point	The point in time within the base time range.

Parameter	Description
Custom Database and Table	You can add custom databases and tables as required. - The system databases cannot be restored. Therefore, the database name cannot be **admin**, **local**, or **config**. - The database name cannot contain spaces and the following special characters: ".$\/*?~#:| - A table name cannot use "system" as the prefix. - The length of *<Database name>.<Table name>* cannot exceed 120 characters. The length of *<Database name>.<Table name>.<Index name>* cannot exceed 128 characters, or the restoration may fail. - Ensure that the name of the restored table is different from that of the existing table. Otherwise, the restoration may fail. - After a successful restoration, a new table named ***Original table name*_bak_*Timestamp*** is generated in the instance by default. If the table contains an index, the namespace of the index is changed to ***Original database name.Original table name*_bak_*Timestamp***. You can rename the table later as required. To distinguish the point in time of the custom databases and tables from those synchronized on the right, set the point in time to a different value. The system restores data to the custom databases and tables based on the time configured here.
Type	You can restore data to a database or table. If you perform a table-level restoration and the table does not exist at the required point in time, an empty table is automatically created. If you perform a database-level restore, data will be restored to the database separately, and the table will not be created.

Figure 9-26 Selecting database and table

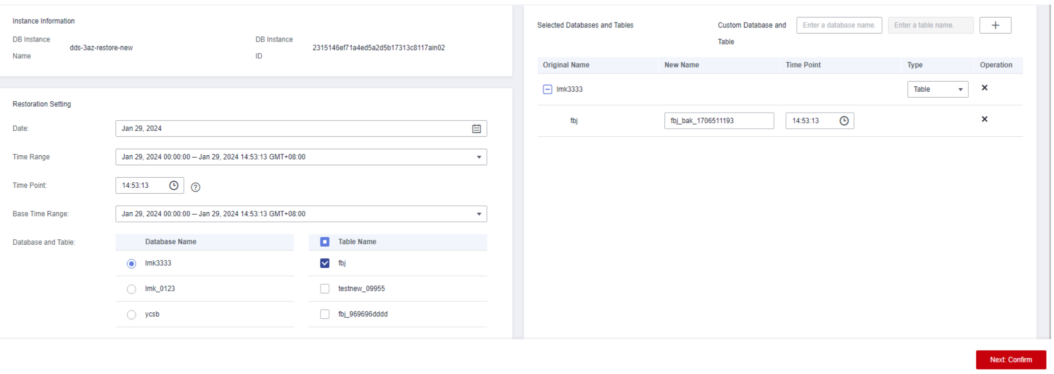

Instance Information

DB Instance Name: dds-3az-restore-new

DB Instance ID: 2315146ef71a4ed5a2d5b17313c8117ain02

Restoration Setting

Date: Jan 29, 2024

Time Range: Jan 29, 2024 00:00:00 — Jan 29, 2024 14:53:13 GMT+08:00

Time Point: 14:53:13

Base Time Range: Jan 29, 2024 00:00:00 — Jan 29, 2024 14:53:13 GMT+08:00

Database and Table:

Database Name	■ Table Name
● lmk3333	☑ fbj
○ lmk_0123	□ testnew_09955
○ ycsb	□ fbj_969696dddd

Selected Databases and Tables

Custom Database and Table: Enter a database name. Enter a table name. +

Original Name	New Name	Time Point	Type	Operation
⊟ lmk3333			Table	×
fbj	fbj_bak_1706511193	14:53:13		×

Step 8 Click **Next: Confirm**.

Step 9 Click **Submit** to start the restoration.

Figure 9-27 Confirming the information

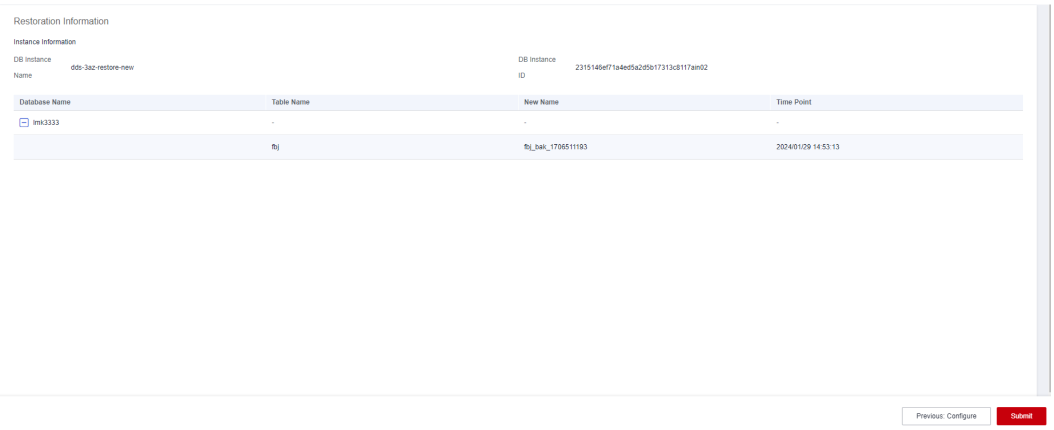

Restoration Information

Instance Information

DB Instance Name: dds-3az-restore-new

DB Instance ID: 2315146ef71a4ed5a2d5b17313c8117ain02

Database Name	Table Name	New Name	Time Point
⊟ lmk3333	-	-	-
	fbj	fbj_bak_1706511193	2024/01/29 14:53:13

Step 10 On the **Instances** page, the DB instance status is **Restoring**. During the restoration process, services are not interrupted.

Step 11 After the restoration is successful, manage data in the database and table as required.

If you need to use the original database and table names, you can use a rename operation to back up the original database and table and switch your service to the restored database and table. Then, delete the original database and table after ensuring that your services are normal.

Example:

db.adminCommand({renameCollection: "db1.test1", to: "db2.test2"})

The above command is used to move the **test1** table from the **db1** database to the **db2** database and rename the table to **test2**.

----End

9.4.4 Restoring a Cross-Region Backup to a Point in Time

DDS allows you to restore cross-region backups to a point in time.

When you enter the point in time that you want to restore the instance to, DDS downloads the most recent full backup file from OBS to the instance. Then, incremental backups are also restored to the specified point in time on the instance. Data is restored at an average speed of 30 MB/s.

Precautions

- To restore backup files to a new instance, your account balance must be greater than or equal to $0 USD. You will pay for the new instance specifications.
- Only cluster and replica set instances of version 4.0 or later can be restored to a point in time.

- Cluster instances can be restored to a specific point in time only after the automated and incremental backup policies are enabled.
- Replica set instances can be restored to a specific point in time only after the automated backup policy is enabled.
- The local database is not included in the databases that can be restored to a specified time point.
- To ensure data security, the dropDatabase operation is blocked when the incremental backup is restored to a point in time. Empty databases or views may exist after the restoration. You can delete them.
- Only whitelisted users can use this function. You need to submit a service ticket to apply for this function. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, choose **Backups** on the navigation pane. Click the **Cross-Region Backups** tab. On the displayed page, locate the target DB instance and click **View Cross-Region Backup** in the **Operation** column.

Figure 9-28 Cross-region backups

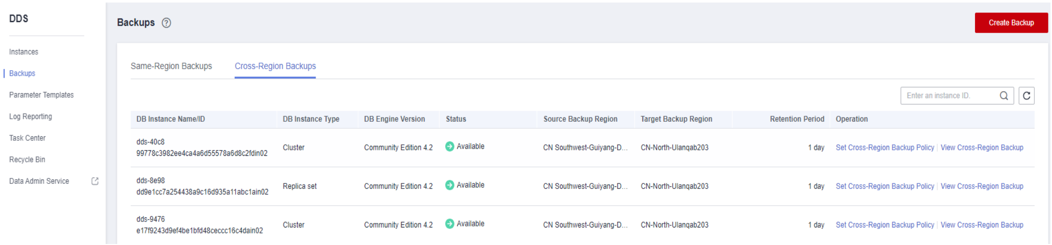

DB Instance Name/ID	DB Instance Type	DB Engine Version	Status	Source Backup Region	Target Backup Region	Retention Period	Operation
dds-40c8 99778c3982ee4ca4a6d55578a6d8c2fdin02	Cluster	Community Edition 4.2	Available	CN Southwest-Guiyang-D...	CN-North-Ulanqab203	1 day	Set Cross-Region Backup Policy | View Cross-Region Backup
dds-8e98 dd9e1cc7a254438a9c16d935a11abc1ain02	Replica set	Community Edition 4.2	Available	CN Southwest-Guiyang-D...	CN-North-Ulanqab203	1 day	Set Cross-Region Backup Policy | View Cross-Region Backup
dds-9476 e17f9243d9ef4be1bfd48ceccc16c4dain02	Cluster	Community Edition 4.2	Available	CN Southwest-Guiyang-D...	CN-North-Ulanqab203	1 day	Set Cross-Region Backup Policy | View Cross-Region Backup

Step 5 Locate the backup to be restored and click **Restore to Point in Time**.

Figure 9-29 Restoring a cross-region backup

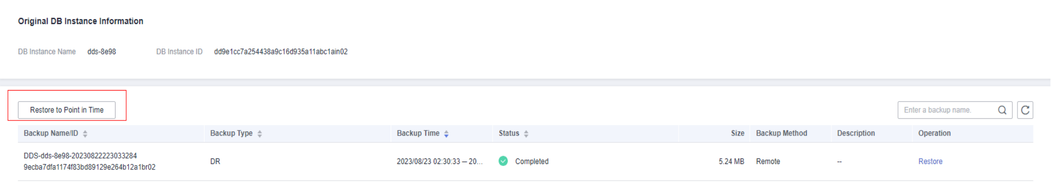

Original DB Instance Information

DB Instance Name dds-8e98 DB Instance ID dd9e1cc7a254438a9c16d935a11abc1ain02

Backup Name/ID	Backup Type	Backup Time	Status	Size	Backup Method	Description	Operation
DDS-dds-8e98-20230822223033284 9ecba7dfa1174f83bd89129e264b12a1br02	DR	2023/08/23 02:30:33 – 20...	Completed	5.24 MB	Remote	--	Restore

Step 6 In the **Restore to Point in Time** dialog box, select **Create New Instance** for **Restoration Method** and click **OK**.

Figure 9-30 Restoring a cross-region backup to a point in time

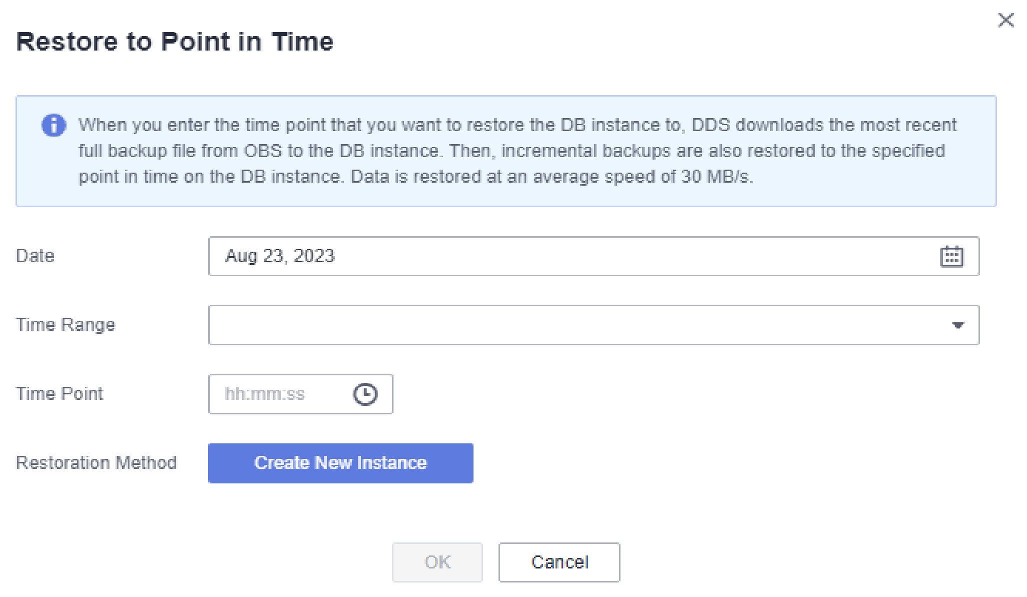

Restore to Point in Time

When you enter the time point that you want to restore the DB instance to, DDS downloads the most recent full backup file from OBS to the DB instance. Then, incremental backups are also restored to the specified point in time on the DB instance. Data is restored at an average speed of 30 MB/s.

Date: Aug 23, 2023

Time Range:

Time Point: hh:mm:ss

Restoration Method: Create New Instance

Step 7 The **Create New Instance** page is displayed for you to create an instance using the backup data. The new DB instance is independent from the original one.

- You are recommended to deploy the restored DB instance in a different AZ to ensure that applications will not be adversely affected by the failure in any single AZ.
- The database type, DB instance type, compatible MongoDB version, storage engine, storage type, and shard quantity must be the same as those of the original and cannot be changed.
- The storage space is the same as that of the original instance by default. You can increase the storage space, but you cannot reduce it.
- Other settings are the same as those of the original DB instance by default and can be modified. For details, see **Buying a Cluster Instance** or **Buying a Replica Set Instance**.
- A full backup is triggered after the new instance is created.

----End

9.5 Restoring Data to an On-Premises Database

9.5.1 Restoring a Cluster Backup to an On-premises Database

9.5.1.1 Overview

This section uses the Linux operating system as an example to describe how to restore the downloaded backup file of a cluster to your on-premises database. For details about how to download backup files, see **Downloading Backup Files**.

Precautions

- This method applies only to cluster instances.
- Only DDS 3.4 and 4.0 instances can be restored in this method. DDS 4.2 or later does not support this method.
- The directories, IP addresses, and ports provided in the example are for reference only. Configure these items based on your service requirements.
- There is one backup file of the configsvr node and multiple backup files of the shardsrv node. The number of backup files depends on the number of shardsvr nodes.
- After the backup file is downloaded, decompress the file using LZ4. Command for reference: **lz4 -d $1 | tar -xC $2**

  $1: indicates the downloaded backup file.

  $2: indicates the directory to which the backup file is decompressed.
- For details about how to migrate data at the database or collection level, see **Migrating Data Using mongodump and mongorestore**.

Prerequisites

MongoDB client 3.4 or 4.0 has been installed on your on-premises database.

9.5.1.2 Directories and Configurations

NOTICE

The local directory, configuration file, and configuration information are not fixed and can be customized.

The following uses backup files of two shardsvr cluster instances as an example (instance ID: cac1efc8e65e42ecad8953352321bfeein02).

- Directory of the decompressed backup files of the configsvr node: /compile/download/backups/cac1efc8e65e42ecad8953352321bfeein02_41c8a32fb10245899708dea453a8c5c9no02
- Directory of the decompressed backup files of the shardsvr1 node:

  /compile/download/backups/cac1efc8e65e42ecad8953352321bfeein02_6cfa6167d4114d7c8cec5b47f9a78dc5no02
- Directory of the decompressed backup files of the shardsvr2 node:

  /compile/download/backups/cac1efc8e65e42ecad8953352321bfeein02_92b196d2401041a7af869a2a3cab7079no02

Data directories and log directories of the three configsvr nodes

/compile/cluster-restore/cfg1/data/db

/compile/cluster-restore/cfg1/log

/compile/cluster-restore/cfg2/data/db

/compile/cluster-restore/cfg2/log

/compile/cluster-restore/cfg3/data/db

/compile/cluster-restore/cfg3/log

Data directories and log directories of the three nodes of shardsvr1

/compile/cluster-restore/shd11/data/db

/compile/cluster-restore/shd11/log

/compile/cluster-restore/shd12/data/db

/compile/cluster-restore/shd12/log

/compile/cluster-restore/shd13/data/db

/compile/cluster-restore/shd13/log

Data directories and log directories of the three nodes of shardsvr2

/compile/cluster-restore/shd21/data/db

/compile/cluster-restore/shd21/log

/compile/cluster-restore/shd22/data/db

/compile/cluster-restore/shd22/log

/compile/cluster-restore/shd23/data/db

/compile/cluster-restore/shd23/log

Log directories of the dds mongos node

/compile/cluster-restore/mgs1/log

/compile/cluster-restore/mgs2/log

IP Address and Port Information

The IP address bound to the process is 127.0.0.1. The port numbers are allocated as follows:

- dds mongos node: 40301, 40302
- configsvr node: 40303, 40304, 40305
- shardsvr1: 40306, 40307, and 40308
- shardsvr2: 40309, 40310, and 40311

Configuration file description

- Configuration file of a single node and configuration files of three nodes in the configsvr replica set

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/single_40303.yaml

/compile/mongodb/mongodb-src-4.0.3/restoreconfig/configsvr_40303.yaml

/compile/mongodb/mongodb-src-4.0.3/restoreconfig/configsvr_40304.yaml

/compile/mongodb/mongodb-src-4.0.3/restoreconfig/configsvr_40305.yaml

- Configuration file of a single node and configuration files of three nodes in the shardsvr1 replica set

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/single_40306.yaml

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/shardsvr_40306.yaml

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/shardsvr_40307.yaml

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/shardsvr_40308.yaml

- Configuration file of a single node and configuration files of three nodes in the shardsvr2 replica set:

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/single_40309.yaml

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/shardsvr_40309.yaml

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/shardsvr_40310.yaml

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/shardsvr_40311.yaml

- Configuration file of the dds mongos node:

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/mongos_40301.yaml

  /compile/mongodb/mongodb-src-4.0.3/restoreconfig/mongos_40302.yaml

Procedure

Command running directory: /compile/mongodb/mongodb-src-4.0.3

9.5.1.3 Restoring the configsvr Replica Set

Preparing Directories

rm -rf /compile/cluster-restore/cfg*

mkdir -p /compile/cluster-restore/cfg1/data/db

mkdir -p /compile/cluster-restore/cfg1/log

mkdir -p /compile/cluster-restore/cfg2/data/db

mkdir -p /compile/cluster-restore/cfg2/log

mkdir -p /compile/cluster-restore/cfg3/data/db

mkdir -p /compile/cluster-restore/cfg3/log

Procedure

Step 1 Prepare the configuration file and data directory of a single node and start the process in single-node mode.

1. The configuration file is as follows (restoreconfig/single_40303.yaml):

```
net:
  bindIp: 127.0.0.1
  port: 40303
  unixDomainSocket: {enabled: false}
```

```
processManagement: {fork: true, pidFilePath: /compile/cluster-restore/cfg1/configsvr.pid}
storage:
  dbPath: /compile/cluster-restore/cfg1/data/db/
  directoryPerDB: true
  engine: wiredTiger
  wiredTiger:
    collectionConfig: {blockCompressor: snappy}
    engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
    indexConfig: {prefixCompression: true}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/cfg1/log/configsingle.log}
```

2. Copy the decompressed **configsvr** file to the **dbPath** directory on the single node.

   ```
   cp -aR
   /compile/download/backups/cac1efc8e65e42ecad8953352321bfeein02_41c8a32fb10245899708dea453a8c5c9no02/* /compile/cluster-restore/cfg1/data/db/
   ```

3. Start the process.

   ```
   ./mongod -f restoreconfig/single_40303.yaml
   ```

Step 2 Connect to the single node and run the following configuration command:

./mongo --host 127.0.0.1 --port 40303

1. Run the following commands to modify the replica set configuration:

   ```
   var cf=db.getSiblingDB('local').system.replset.findOne();
   cf['members'][0]['host']='127.0.0.1:40303';
   cf['members'][1]['host']='127.0.0.1:40304';
   cf['members'][2]['host']='127.0.0.1:40305';
   cf['members'][0]['hidden']=false;
   cf['members'][1]['hidden']=false;
   cf['members'][2]['hidden']=false;
   cf['members'][0]['priority']=1;
   cf['members'][1]['priority']=1;
   cf['members'][2]['priority']=1;
   db.getSiblingDB('local').system.replset.remove({});
   db.getSiblingDB('local').system.replset.insert(cf)
   ```

2. Run the following commands to clear the built-in accounts:

   ```
   db.getSiblingDB('admin').dropAllUsers();
   db.getSiblingDB('admin').dropAllRoles();
   ```

3. Run the following command to update the dds mongos and shard information:

   ```
   db.getSiblingDB('config').mongos.remove({});
   ```

   Query the _id information about multiple shards in the **config.shards** table. The _id information is used as the query condition of _id in the following statements. Update records in sequence.

   ```
   db.getSiblingDB('config').shards.update({'_id' : 'shard_1'},{$set: {'host': 'shard_1/127.0.0.1:40306,127.0.0.1:40307,127.0.0.1:40308'}})
   ```

```
db.getSiblingDB('config').shards.update({'_id' : 'shard_2'},{$set: {'host': 'shard_2/127.0.0.1:40309,127.0.0.1:40310,127.0.0.1:40311'}})
```

```
db.getSiblingDB('config').mongos.find({});
```

```
db.getSiblingDB('config').shards.find({});
```

4. Run the following command to stop the single-node process:

   ```
   db.getSiblingDB('admin').shutdownServer();
   ```

Step 3 Create a configsvr replica set.

1. Copy the **dbPath** file of the configsvr1 node to the directories of the other two configsvr nodes.

   ```
   cp -aR /compile/cluster-restore/cfg1/data/db/ /compile/cluster-restore/cfg2/data/db/
   ```

   ```
   cp -aR /compile/cluster-restore/cfg1/data/db/ /compile/cluster-restore/cfg3/data/db/
   ```

2. Add the replica set configuration attribute to the configuration file (**restoreconfig/configsvr_40303.yaml**) of the configsvr-1 node.

   ```
   net:
     bindIp: 127.0.0.1
     port: 40303
     unixDomainSocket: {enabled: false}
   processManagement: {fork: true, pidFilePath: /compile/cluster-restore/cfg1/configsvr.pid}
   replication: {replSetName: config}
   sharding: {archiveMovedChunks: false, clusterRole: configsvr}
   storage:
     dbPath: /compile/cluster-restore/cfg1/data/db/
     directoryPerDB: true
     engine: wiredTiger
     wiredTiger:
       collectionConfig: {blockCompressor: snappy}
       engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
       indexConfig: {prefixCompression: true}
   systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/cfg1/log/configsvr.log}
   ```

3. Start the process.

   ```
   ./mongod -f restoreconfig/configsvr_40303.yaml
   ```

4. Add the replica set configuration attribute to the configuration file (**restoreconfig/configsvr_40304.yaml**) of the configsvr-2 node.

   ```
   net:
     bindIp: 127.0.0.1
     port: 40304
     unixDomainSocket: {enabled: false}
   processManagement: {fork: true, pidFilePath: /compile/cluster-restore/cfg2/configsvr.pid}
   replication: {replSetName: config}
   sharding: {archiveMovedChunks: false, clusterRole: configsvr}
   storage:
     dbPath: /compile/cluster-restore/cfg2/data/db/
     directoryPerDB: true
     engine: wiredTiger
     wiredTiger:
       collectionConfig: {blockCompressor: snappy}
       engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
       indexConfig: {prefixCompression: true}
   systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/cfg2/log/configsvr.log}
   ```

5. Start the process.

   ```
   ./mongod -f restoreconfig/configsvr_40304.yaml
   ```

6. Add the replica set configuration attribute to the configuration file (**restoreconfig/configsvr_40305.yaml**) of the configsvr-3 node.

   ```
   net:
     bindIp: 127.0.0.1
     port: 40305
     unixDomainSocket: {enabled: false}
   processManagement: {fork: true, pidFilePath: /compile/cluster-restore/cfg3/configsvr.pid}
   replication: {replSetName: config}
   sharding: {archiveMovedChunks: false, clusterRole: configsvr}
   storage:
     dbPath: /compile/cluster-restore/cfg3/data/db/
     directoryPerDB: true
     engine: wiredTiger
     wiredTiger:
       collectionConfig: {blockCompressor: snappy}
       engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
       indexConfig: {prefixCompression: true}
   systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/cfg3/log/configsvr.log}
   ```

7. Start the process.

   ```
   ./mongod -f restoreconfig/configsvr_40305.yaml
   ```

Step 4 Wait until the primary node is selected.

```
./mongo --host 127.0.0.1 --port 40303
```

Run the **rs.status()** command to check whether the primary node exists.

----End

9.5.1.4 Restoring the shardsvr1 Replica Set

Preparing Directories

```
rm -rf /compile/cluster-restore/shd1*
mkdir -p /compile/cluster-restore/shd11/data/db
mkdir -p /compile/cluster-restore/shd11/log
mkdir -p /compile/cluster-restore/shd12/data/db
mkdir -p /compile/cluster-restore/shd12/log
mkdir -p /compile/cluster-restore/shd13/data/db
mkdir -p /compile/cluster-restore/shd13/log
```

Procedure

Step 1 Prepare the configuration file and directory of a single node and start the process in single-node mode.

1. The configuration file is as follows (**restoreconfig/single_40306.yaml**):

   ```
   net:
     bindIp: 127.0.0.1
   ```

```
  port: 40306
  unixDomainSocket: {enabled: false}
processManagement: {fork: true, pidFilePath: /compile/cluster-restore/shd11/mongod.pid}
storage:
  dbPath: /compile/cluster-restore/shd11/data/db/
  directoryPerDB: true
  engine: wiredTiger
  wiredTiger:
    collectionConfig: {blockCompressor: snappy}
    engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
    indexConfig: {prefixCompression: true}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/shd11/log/mongod.log}
```

2. Copy the decompressed **shardsvr1** file to the **dbPath** directory on the single node.

   ```
   cp -aR
   /compile/download/backups/cac1efc8e65e42ecad8953352321bfeein02_6cfa6167d4114d7c8cec5b47f9a78dc5no02/* /compile/cluster-restore/shd11/data/db/
   ```

3. Start the process.

   ```
   ./mongod -f restoreconfig/single_40306.yaml
   ```

Step 2 Connect to the single node and run the following configuration command:

Connection command: ./mongo --host 127.0.0.1 --port 40306

1. Run the following commands to modify the replica set configuration:

   ```
   var cf=db.getSiblingDB('local').system.replset.findOne();
   cf['members'][0]['host']='127.0.0.1:40306';
   cf['members'][1]['host']='127.0.0.1:40307';
   cf['members'][2]['host']='127.0.0.1:40308';
   cf['members'][0]['hidden']=false;
   cf['members'][1]['hidden']=false;
   cf['members'][2]['hidden']=false;
   cf['members'][0]['priority']=1;
   cf['members'][1]['priority']=1;
   cf['members'][2]['priority']=1;
   db.getSiblingDB('local').system.replset.remove({});
   db.getSiblingDB('local').system.replset.insert(cf)
   ```

2. Run the following commands to clear the built-in accounts:

   ```
   db.getSiblingDB('admin').dropAllUsers();
   db.getSiblingDB('admin').dropAllRoles();
   ```

3. Run the following commands to update the configsvr information:

   Connection command: ./mongo --host 127.0.0.1 --port 40306

   ```
   var vs = db.getSiblingDB('admin').system.version.find();
   while (vs.hasNext()) {
   var curr = vs.next();
   if (curr.hasOwnProperty('configsvrConnectionString')) {
   ```

```
db.getSiblingDB('admin').system.version.update({'_id' : curr._id}, {$set: {'configsvrConnectionString': 'config/127.0.0.1:40303,127.0.0.1:40304,127.0.0.1:40305'}});
}
}
```

4. Run the following command to stop the single-node process:

   ```
   db.getSiblingDB('admin').shutdownServer();
   ```

Step 3 Create the shardsvr1 replica set.

1. Copy the **dbPath** file of the shardsvr1 node to the directories of the other two shardsvr nodes.

   ```
   cp -aR /compile/cluster-restore/shd11/data/db/ /compile/cluster-restore/shd12/data/db/
   ```

   ```
   cp -aR /compile/cluster-restore/shd11/data/db/ /compile/cluster-restore/shd13/data/db/
   ```

2. Add the replica set configuration attribute to the configuration file (**restoreconfig/shardsvr_40306.yaml**) of the shardsvr1-1 node.

   --- For details about the value of **replication.replSetName**, see the shard _id information in **Step 2.3**.

   ```
   net:
     bindIp: 127.0.0.1
     port: 40306
     unixDomainSocket: {enabled: false}
   processManagement: {fork: true, pidFilePath: /compile/cluster-restore/shd11/mongod.pid}
   replication: {replSetName: shard_1}
   sharding: {archiveMovedChunks: false, clusterRole: shardsvr}
   storage:
     dbPath: /compile/cluster-restore/shd11/data/db/
     directoryPerDB: true
     engine: wiredTiger
     wiredTiger:
       collectionConfig: {blockCompressor: snappy}
       engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
       indexConfig: {prefixCompression: true}
   systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/shd11/log/mongod.log}
   ```

3. Start the process.

   ```
   ./mongod -f restoreconfig/shardsvr_40306.yaml
   ```

4. Add the replica set configuration attribute to the configuration file (**restoreconfig/shardsvr_40307.yaml**) of the shardsvr1-2 node.

   --- For details about the value of **replication.replSetName**, see the shard _id information in **Step 2.3**.

   ```
   net:
     bindIp: 127.0.0.1
     port: 40307
     unixDomainSocket: {enabled: false}
   processManagement: {fork: true, pidFilePath: /compile/cluster-restore/shd12/mongod.pid}
   replication: {replSetName: shard_1}
   sharding: {archiveMovedChunks: false, clusterRole: shardsvr}
   storage:
     dbPath: /compile/cluster-restore/shd12/data/db/
     directoryPerDB: true
     engine: wiredTiger
   ```

```
  wiredTiger:
    collectionConfig: {blockCompressor: snappy}
    engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
    indexConfig: {prefixCompression: true}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/shd12/log/mongod.log}
```

5. Start the process.

   ./mongod -f restoreconfig/shardsvr_40307.yaml

6. Add the replica set configuration attribute to the configuration file (**restoreconfig/shardsvr_40308.yaml**) of the shardsvr1-3 node.

   --- For details about the value of **replication.replSetName**, see the shard _id information in **Step 2.3**.

```
net:
  bindIp: 127.0.0.1
  port: 40308
  unixDomainSocket: {enabled: false}
processManagement: {fork: true, pidFilePath: /compile/cluster-restore/shd13/mongod.pid}
replication: {replSetName: shard_1}
sharding: {archiveMovedChunks: false, clusterRole: shardsvr}
storage:
  dbPath: /compile/cluster-restore/shd13/data/db/
  directoryPerDB: true
  engine: wiredTiger
  wiredTiger:
    collectionConfig: {blockCompressor: snappy}
    engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
    indexConfig: {prefixCompression: true}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/shd13/log/mongod.log}
```

7. Start the process.

   ./mongod -f restoreconfig/shardsvr_40308.yaml

Step 4 Wait until the primary node is selected.

./mongo --host 127.0.0.1 --port 40306

Run the **rs.status()** command to check whether the primary node exists.

----End

9.5.1.5 Restoring the shardsvr2 Replica Set

Preparing Directories

rm -rf /compile/cluster-restore/shd2*

mkdir -p /compile/cluster-restore/shd21/data/db

mkdir -p /compile/cluster-restore/shd21/log

mkdir -p /compile/cluster-restore/shd22/data/db

mkdir -p /compile/cluster-restore/shd22/log

mkdir -p /compile/cluster-restore/shd23/data/db

mkdir -p /compile/cluster-restore/shd23/log

Procedure

Step 1 Prepare the configuration file and directory of a single node and start the process in single-node mode.

1. The configuration file is as follows (**restoreconfig/single_40309.yaml**):

```
net:
  bindIp: 127.0.0.1
  port: 40309
  unixDomainSocket: {enabled: false}
processManagement: {fork: true, pidFilePath: /compile/cluster-restore/shd21/mongod.pid}
storage:
  dbPath: /compile/cluster-restore/shd21/data/db/
  directoryPerDB: true
  engine: wiredTiger
  wiredTiger:
    collectionConfig: {blockCompressor: snappy}
    engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
    indexConfig: {prefixCompression: true}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/shd21/log/mongod.log}
```

1. Copy the decompressed **shardsvr2** file to the **dbPath** directory on the single node.

```
cp -aR
/compile/download/backups/cac1efc8e65e42ecad8953352321bfeein02_92b196d2401041a7af869a2a3cab7079no02/* /compile/cluster-restore/shd21/data/db/
```

2. Start the process.

```
./mongod -f restoreconfig/single_40309.yaml
```

Step 2 Connect to the single node and run the following configuration command:

Connection command: ./mongo --host 127.0.0.1 --port 40309

1. Run the following commands to modify the replica set configuration:

```
var cf=db.getSiblingDB('local').system.replset.findOne();
cf['members'][0]['host']='127.0.0.1:40309';
cf['members'][1]['host']='127.0.0.1:40310';
cf['members'][2]['host']='127.0.0.1:40311';
cf['members'][0]['hidden']=false;
cf['members'][1]['hidden']=false;
cf['members'][2]['hidden']=false;
cf['members'][0]['priority']=1;
cf['members'][1]['priority']=1;
cf['members'][2]['priority']=1;
db.getSiblingDB('local').system.replset.remove({});
db.getSiblingDB('local').system.replset.insert(cf)
```

2. Run the following commands to clear the built-in accounts:

```
db.getSiblingDB('admin').dropAllUsers();
db.getSiblingDB('admin').dropAllRoles();
```

3. Run the following commands to update the configsvr information:

   ```
   var vs = db.getSiblingDB('admin').system.version.find();
   while (vs.hasNext()) {
   var curr = vs.next();
   if (curr.hasOwnProperty('configsvrConnectionString')) {
   db.getSiblingDB('admin').system.version.update({'_id' : curr._id}, {$set: {'configsvrConnectionString': 'config/127.0.0.1:40303,127.0.0.1:40304,127.0.0.1:40305'}});
   }
   }
   ```

4. Run the following command to stop the single-node process:

   ```
   db.getSiblingDB('admin').shutdownServer();
   ```

Step 3 Create the shardsvr2 replica set.

1. Copy the **dbPath** file of the shardsvr2 node to the directories of the other two shardsvr nodes.

   ```
   cp -aR /compile/cluster-restore/shd21/data/db/ /compile/cluster-restore/shd22/data/db/
   ```

   ```
   cp -aR /compile/cluster-restore/shd21/data/db/ /compile/cluster-restore/shd23/data/db/
   ```

2. Add the replica set configuration attribute to the configuration file (**restoreconfig/shardsvr_40309.yaml**) of the shardsvr2-1 node.

   --- For details about the value of **replication.replSetName**, see the shard _id information in **Step 2.3**.

   ```
   net:
     bindIp: 127.0.0.1
     port: 40309
     unixDomainSocket: {enabled: false}
   processManagement: {fork: true, pidFilePath: /compile/cluster-restore/shd21/mongod.pid}
   replication: {replSetName: shard_2}
   sharding: {archiveMovedChunks: false, clusterRole: shardsvr}
   storage:
     dbPath: /compile/cluster-restore/shd21/data/db/
     directoryPerDB: true
     engine: wiredTiger
     wiredTiger:
       collectionConfig: {blockCompressor: snappy}
       engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
       indexConfig: {prefixCompression: true}
   systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/shd21/log/mongod.log}
   ```

3. Start the process.

   ```
   ./mongod -f restoreconfig/shardsvr_40309.yaml
   ```

4. Add the replica set configuration attribute to the configuration file (**restoreconfig/shardsvr_40310.yaml**) of the shardsvr2-2 node.

   --- For details about the value of **replication.replSetName**, see the shard _id information in **Step 2.3**.

   ```
   net:
     bindIp: 127.0.0.1
     port: 40310
   ```

```
  unixDomainSocket: {enabled: false}
processManagement: {fork: true, pidFilePath: /compile/cluster-restore/shd22/mongod.pid}
replication: {replSetName: shard_2}
sharding: {archiveMovedChunks: false, clusterRole: shardsvr}
storage:
  dbPath: /compile/cluster-restore/shd22/data/db/
  directoryPerDB: true
  engine: wiredTiger
  wiredTiger:
    collectionConfig: {blockCompressor: snappy}
    engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
    indexConfig: {prefixCompression: true}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/shd22/log/mongod.log}
```

5. Start the process.

   ./mongod -f restoreconfig/shardsvr_40310.yaml

6. Add the replica set configuration attribute to the configuration file (**restoreconfig/shardsvr_40311.yaml**) of the shardsvr2-2 node.

   --- For details about the value of **replication.replSetName**, see the shard _id information in **Step 2.3**.

```
net:
  bindIp: 127.0.0.1
  port: 40311
  unixDomainSocket: {enabled: false}
processManagement: {fork: true, pidFilePath: /compile/cluster-restore/shd23/mongod.pid}
replication: {replSetName: shard_2}
sharding: {archiveMovedChunks: false, clusterRole: shardsvr}
storage:
  dbPath: /compile/cluster-restore/shd23/data/db/
  directoryPerDB: true
  engine: wiredTiger
  wiredTiger:
    collectionConfig: {blockCompressor: snappy}
    engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
    indexConfig: {prefixCompression: true}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/shd23/log/mongod.log}
```

7. Start the process.

   ./mongod -f restoreconfig/shardsvr_40311.yaml

Step 4 Wait until the primary node is selected.

./mongo --host 127.0.0.1 --port 40309

Run the **rs.status()** command to check whether the primary node exists.

----End

9.5.1.6 Restoring the dds mongos Node

Step 1 Prepare the configuration file and directory of the dds mongos node.

rm -rf /compile/cluster-restore/mgs*

mkdir -p /compile/cluster-restore/mgs1/log

mkdir -p /compile/cluster-restore/mgs2/log

Step 2 Configuration file (restoreconfig/mongos_40301.yaml)

```
net:
  bindIp: 127.0.0.1
  port: 40301
  unixDomainSocket: {enabled: false}
processManagement: {fork: true, pidFilePath: /compile/cluster-restore/mgs1/mongos.pid}
sharding: {configDB: 'config/127.0.0.1:40303,127.0.0.1:40304,127.0.0.1:40305'}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/
mgs1/log/mongos.log}
```

Step 3 Configuration file (restoreconfig/mongos_40302.yaml)

```
net:
  bindIp: 127.0.0.1
  port: 40302
  unixDomainSocket: {enabled: false}
processManagement: {fork: true, pidFilePath: /compile/cluster-restore/mgs2/mongos.pid}
sharding: {configDB: 'config/127.0.0.1:40303,127.0.0.1:40304,127.0.0.1:40305'}
systemLog: {destination: file, logAppend: true, logRotate: reopen, path: /compile/cluster-restore/
mgs2/log/mongos.log}
```

Step 4 Start the mongo node.

./mongos -f restoreconfig/mongos_40301.yaml

./mongos -f restoreconfig/mongos_40302.yaml

----End

9.5.1.7 Checking the Cluster Status

Connect to the cluster through dds mongos and check the data status.

./mongo --host 127.0.0.1 --port 40301

./mongo --host 127.0.0.1 --port 40302

9.5.2 Restoring a Replica Set Backup to an On-Premises Database

To restore a DB instance backup file to an on-premises database, you can only use databases on Linux.

This section uses the Linux operating system as an example to describe how to restore the downloaded backup file of a replica set instance to your on-premises databases. For details about how to download backup files, see **Downloading Backup Files**.

Precautions

- MongoDB client 3.4 has been installed on your on-premises MongoDB database.
- Only DDS 3.4 and 4.0 instances can be restored in this method. DDS 4.2 or later does not support this method.
- For details about how to migrate data at the database or collection level, see **Migrating Data Using mongodump and mongorestore**.

Procedure

Step 1 Log in to the server on which the on-premises databases are deployed.

Assume that **/path/to/mongo** is the directory for restoration, and **/path/to/mongo/data** is the directory for storing the backup file.

Step 2 Before the restoration, ensure that the **/path/to/mongo/data** directory is empty.

cd /path/to/mongo/data/

rm -rf *

Step 3 Copy and paste the downloaded backup file package to **/path/to/mongo/data/** and decompress it.

lz4 -d xxx_.tar.gz |tar -xC /path/to/mongo/data/

Step 4 Create the **mongod.conf** configuration file in **/path/to/mongo**.

touch mongod.conf

Step 5 Start the database in single-node mode.

1. Modify the **mongod.conf** file to meet the backup startup configuration requirements.

   The following is a configuration template for backup startup:

```
systemLog:
    destination: file
    path: /path/to/mongo/mongod.log
    logAppend: true
security:
    authorization: enabled
storage:
    dbPath: /path/to/mongo/data
    directoryPerDB: true
    engine: wiredTiger
    wiredTiger:
        collectionConfig: {blockCompressor: snappy}
        engineConfig: {directoryForIndexes: true, journalCompressor: snappy}
        indexConfig: {prefixCompression: true}
net:
    http:
        enabled: false
    port: 27017
    bindIp: xxx.xxx.xxx.xxx,xxx.xxx.xxx.xxx
    unixDomainSocket:
        enabled: false
processManagement:
    fork: true
    pidFilePath: /path/to/mongo/mongod.pid
```

   NOTE

   bindIp indicates the IP address bound to the database. This field is optional. If it is not specified, your local IP address is bound by default.

2. Run the **mongod.conf** command to start the database.

   */usr/bin/***mongod -f /path/to/mongo/mongod.conf**

NOTE

/usr/bin/ is the directory that stores the **mongod** file of the installed MongoDB client.

3. After the database is started, log in to the database using mongo shell to verify the restoration result.

   mongo --host *<DB_HOST>* **-u** *<DB_USER>* **-p** *<PASSWORD>* **--authenticationDatabase admin**

   NOTE

   - **DB_HOST** is the IP address bound to the database.
   - **DB_USER** is the database user. The default value is **rwuser**.
   - **PASSWORD** is the password for the database user, which is the password used for backing up the DB instance.

----End

Starting the Database in Replica Set Mode

By default, the physical backup of the DDS DB instance contains the replica set configuration of the original DB instance. You need to start the database in single-node mode. Otherwise, the database cannot be accessed.

If you want to start the database in replica set mode, perform step **Step 5** and then perform the following steps:

Step 1 Log in to the database using mongo shell.

Step 2 Remove the original replica set configuration.

use local

db.system.replset.remove({})

Step 3 Stop the database process.

use admin

db.shutdownServer()

Step 4 Add the replication configuration in the **mongod.conf** file in the **/path/to/mongo/** directory. For details about the command usage, see **Deploy a Replica Set**.

Step 5 Run the **mongod.conf** command to start the database.

*/usr/bin/***mongod -f /path/to/mongo/mongod.conf**

NOTE

/usr/bin/ is the directory that stores the **mongod** file of the installed MongoDB client.

Step 6 Add the replica set members and initialize the replica set.

NOTE

Use the rs.initiate() command to perform the preceding step. For details, see **rs.initiate()**.

----End

9.5.3 Restoring a Single Node Backup to an On-Premises Database

This section uses the Linux operating system as an example to describe how to restore the downloaded backup file of a single node instance to your on-premises database. For details about how to download backup files, see **Downloading Backup Files**.

NOTE

Huawei Cloud has discontinued the sale of DDS single node instances since July 15, 2023.

Precautions

- MongoDB client 3.4 has been installed on your on-premises MongoDB database.
- Only DDS 3.4 and 4.0 instances can be restored in this method. DDS 4.2 or later does not support this method.
- For details about how to migrate data at the database or collection level, see **Migrating Data Using mongodump and mongorestore**.

Procedure

Step 1 **Download the backup file of the single node.**

Step 2 Log in to the device that can access the on-premises database.

Step 3 Upload the single-node backup file to the device that can access the on-premises database.

Select an uploading method based on the OS you are using. In Linux, for example, run the following command:

scp -r *<IDENTITY_DIR>* *<REMOTE_USER>*@*<REMOTE_ADDRESS>*:*<REMOTE_DIR>*

- **IDENTITY_DIR** is the directory that stores the backup file.
- **REMOTE_USER** is the username for logging in to the device that can access the on-premises database.
- **REMOTE_ADDRESS** is the IP address of the host that can access the on-premises database.
- **REMOTE_DIR** is the destination directory to which the backup file is imported.

In Windows, upload the backup file using file transfer tools.

Step 4 Import the backup files in the on-premises database.

./mongorestore --host *<DB_HOST>* **--port** *<DB_PORT>* **-u** *<DB_USER>* **--authenticationDatabase** *<AUTH_DB>* **--drop --gzip --archive=***<Backup directory>* **-vvvv --stopOnError**

- **DB_HOST** is the on-premises database address.
- **DB_PORT** is the on-premises database port.
- **DB_USER** is the on-premises database username.

- **AUTH_DB** is the database that authenticates DB_USER. Generally, this value is **admin**.
- **Backup directory** is the backup file name.

Enter the password for logging in to the on-premises database when prompted:

```
Enter password:
```

Example:

./mongorestore --host 192.168.6.187 --port 8635 -u rwuser --authenticationDatabase admin --drop --gzip --archive=xxx_tar.gz -vvvv --stopOnError

----End

9.6 Restoring Data of Enhanced Edition

Scenarios

DDS 4.4 and DDS 4.4 Enhanced Edition (DDS 4.4 pro for short) are incompatible due to different underlying data storage structures. You can upgrade DDS 4.4 to DDS 4.4 pro in either of the following ways:

- Use DRS to migrate data from DDS 4.4 to DDS 4.4 pro. For details, see **From On-Premises MongoDB to DDS**.
- Contact customer service to upgrade DDS 4.4 to DDS 4.4 pro by rebuilding data on the original DB instance.

Precautions

- After data in a DDS DB instance is migrated to an enhanced binary system, backups before the migration cannot be restored to the original DB instance or a new DB instance.
- After data in a DDS DB instance is migrated to an enhanced binary system, backups cannot be restored to the time point before the migration.
- DDS enhanced binary data backups cannot be restored to on-premises databases.

10 Parameter Template Management

10.1 Overview

DB parameter templates act as a container for engine configuration values that are applied to one or more DB instances. You can customize the parameter settings to manage DB engine configurations.

Parameter Template Type

When creating a DB instance, you can associate a default parameter template or a customized parameter template with the DB instance. After a DB instance is created, you can also change the associated parameter template.

- Default parameter template

  The DB engine parameter values and system service parameter values in the default parameter group are designed for optimizing the database performance.

- Custom parameter template

  If you need a DB instance with customized parameter settings, you can create a parameter template and change the parameter values as required.

  If you change the parameter values of the parameter template associated with several DB instances, the changes will apply to all these DB instances.

Application Scenarios

- If you want to use a customized parameter template, you only need to create a parameter template in advance and select the parameter template when creating a DB instance. For details about how to create a parameter template, see **Creating a Parameter Template**.
- When you have already created a parameter template and want to include most of the custom parameters and values from that template in a new parameter template, you can replicate that parameter template following the instructions provided in section **Replicating a Parameter Template**.

Precautions

- Default parameter templates are unchangeable. You can only view them by clicking their names. If inappropriate settings of customized parameter templates lead to a database startup failure, you can reset the customized parameter template by referring to the settings of the default parameter template.
- After modifying a parameter, you need to view the associated instance status in the instance list. If **Pending restart** is displayed, you need to restart the instance for the modification to take effect.
- Improperly setting parameters in a parameter template may have unintended adverse effects, including degraded performance and system instability. Exercise caution when modifying database parameters and you need to back up data before modifying parameters in a parameter template. Before applying parameter changes to a production DB instance, you should try out these changes on a test DB instance.

10.2 Creating a Parameter Template

DB parameter templates act as a container for engine configuration values that are applied to one or more DB instances.

Precautions

- DDS does not share parameter template quotas with RDS.
- Each account can create up to 100 DDS parameter templates for the cluster, replica set, and single node instances.

Cluster

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Create Parameter Template**.

Step 6 Select **Cluster** for **DB Instance Type**, specify **DB Engine Version**, **Node Type**, **New Parameter Template**, and **Description** (optional), and then click **OK**.

- **Node Type**: specifies the node type that this parameter template will apply to. For example, to create a parameter template applying to config, select **config**.
- **New Parameter Template**: The template name can be up to 64 characters. It must start with a letter and can contain only letters (case-sensitive), digits, hyphens (-), periods (.), and underscores (_).
- **Description**: It can contain up to 256 characters but cannot contain line breaks or the following special characters >!<"&'=

Step 7 On the **Parameter Templates** page, view and manage parameter templates on the **Clusters** tab.

----End

Replica Set

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Create Parameter Template**.

Step 6 Select **Replica set** for **DB Instance Type**, specify **DB Engine Version**, **Node Type**, **Parameter Template Name**, and **Description** (optional), and then click **OK**.

- **Node Type**: specifies the node type that this parameter template will apply to. For example, to create a parameter template applying to a read replica, select **readonly**.
- **New Parameter Template**: The template name can be up to 64 characters. It must start with a letter and can contain only letters (case-sensitive), digits, hyphens (-), periods (.), and underscores (_).
- **Description**: It can contain up to 256 characters but cannot contain line breaks or the following special characters >!<"&'=

Step 7 On the **Parameter Templates** page, view and manage parameter templates on the **Replica Sets** tab.

----End

Single Node

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Create Parameter Template**.

Step 6 Select **Single node** for **DB Instance Type**, specify **DB Engine Version**, **New Parameter Template**, and **Description** (optional), and then click **OK**.

- **New Parameter Template**: The template name can be up to 64 characters. It must start with a letter and can contain only letters (case-sensitive), digits, hyphens (-), periods (.), and underscores (_).
- **Description**: It can contain up to 256 characters but cannot contain line breaks or the following special characters >!<"&'=

Step 7 On the **Parameter Templates** page, view and manage parameter templates on the **Single Nodes** tab.

----End

10.3 Modifying DDS DB Instance Parameters

You can modify parameters in custom parameter templates as needed to enjoy better performance of DDS.

You can modify parameters in either of the following ways:

- Directly modify the parameters of a specified instance.

  If you change a dynamic parameter value in a parameter template and save the change, the change takes effect immediately regardless of the **Effective upon Reboot** setting. If you modify static parameters on the **Parameters** page of an instance and save the modifications, the modifications take effect only after you manually restart the target instance.
- Modify the parameters in a parameter template and apply the template to the instance.

  The changes only take effect after you apply the template to the instance. If you modify static parameters in a custom parameter template on the **Parameter Template Management** page and save the modifications, the modifications take effect only after you apply the parameter template to instances and manually restart the instances. For details about how to apply a parameter template to instances, see **Applying a Parameter Template**.

Precautions

- You can change parameter values in custom parameter templates but cannot change the default parameter templates provided by the system. You can only click the name of a default parameter template to view its details.
- If a custom parameter template is set incorrectly, the instance associated with the template may fail to start. You can re-configure the custom parameter template according to the configurations of the default parameter template.

⚠ CAUTION

Exercise caution when modifying parameter values to prevent exceptions.

Modifying Parameters of an Instance

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Instances**. On the displayed page, click the DB instance whose parameters you wish to modify.

Step 5 In the navigation pane on the left, choose **Parameters**. On the displayed page, modify parameters as required.

Figure 10-1 Modifying parameters of an instance

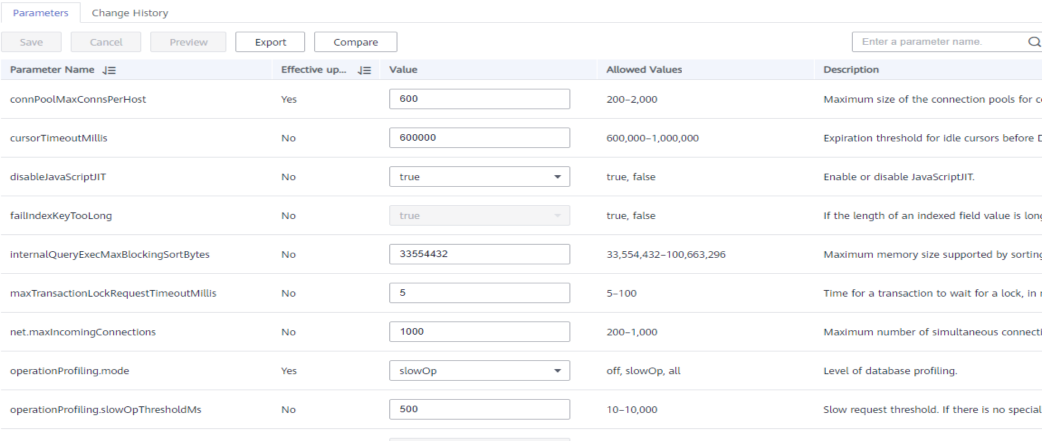

Parameter Name	Effective up...	Value	Allowed Values	Description
connPoolMaxConnsPerHost	Yes	600	200–2,000	Maximum size of the connection pools for c
cursorTimeoutMillis	No	600000	600,000–1,000,000	Expiration threshold for idle cursors before D
disableJavaScriptJIT	No	true	true, false	Enable or disable JavaScriptJIT.
failIndexKeyTooLong	No	true	true, false	If the length of an indexed field value is lon
internalQueryExecMaxBlockingSortBytes	No	33554432	33,554,432–100,663,296	Maximum memory size supported by sortin
maxTransactionLockRequestTimeoutMillis	No	5	5–100	Time for a transaction to wait for a lock, in
net.maxIncomingConnections	No	1000	200–1,000	Maximum number of simultaneous connecti
operationProfiling.mode	Yes	slowOp	off, slowOp, all	Level of database profiling.
operationProfiling.slowOpThresholdMs	No	500	10–10,000	Slow request threshold. If there is no special

Step 6 Modify parameters based on the DB instance type.

- If the DB instance is a cluster instance, select dds mongos, shard, or config on the **Parameters** page and change the value of **net.maxIncomingConnections**, which indicates maximum number of concurrent connections that dds mongos or mongod can be connected.

  Enter **net.maxIncomingConnections** in the search box in the upper right corner of the page and click the search icon to search for this parameter.

  Figure 10-2 Changing the maximum number of connections

  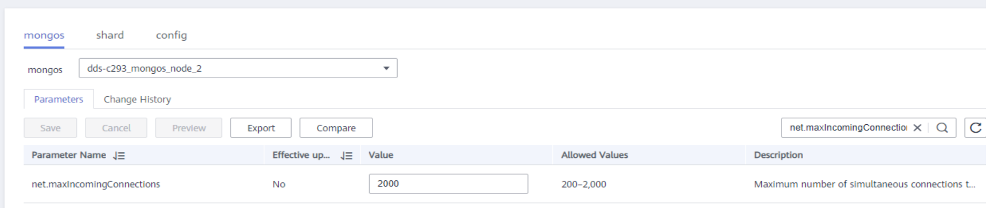

  mongos: dds-c293_mongos_node_2

  net.maxIncomingConnectio

Parameter Name	Effective up...	Value	Allowed Values	Description
net.maxIncomingConnections	No	2000	200–2,000	Maximum number of simultaneous connections t...

- If the DB instance is a replica set instance, select **Replica set nodes** or **Read replicas** on the **Parameters** page and change the value of **net.maxIncomingConnections**, which indicates maximum number of concurrent connections that dds mongos or mongod can be connected.

  Enter **net.maxIncomingConnections** in the search box in the upper right corner of the page and click the search icon to search for this parameter.

  Figure 10-3 Changing the maximum number of connections

  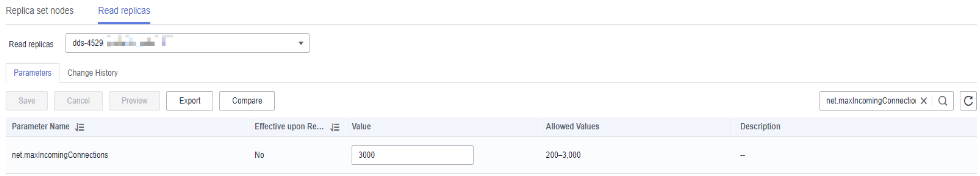

  Read replicas: dds-4529 [redacted]

  net.maxIncomingConnectio

Parameter Name	Effective upon Re...	Value	Allowed Values	Description
net.maxIncomingConnections	No	3000	200–3,000	--

- If the DB instance is a single node instance, change the value of **net.maxIncomingConnections**, which indicates maximum number of concurrent connections that dds mongos or mongod can be connected.

  Enter **net.maxIncomingConnections** in the search box in the upper right corner of the page and click the search icon to search for this parameter.

Figure 10-4 Changing the maximum number of connections

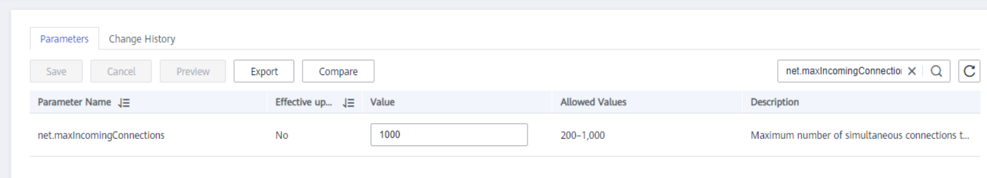

Parameter Name	Effective up...	Value	Allowed Values	Description
net.maxIncomingConnections	No	1000	200–1,000	Maximum number of simultaneous connections t...

Step 7 Change the maximum number of connections based on the parameter value range and instance specifications. This default value depends on the DB instance specifications. This parameter is displayed as **default** before being set, indicating that the parameter value varies with the memory specifications. For details about the parameters, see **Parameters**.

- To save the changes, click **Save**.
- If you want to cancel the modifications, click **Cancel**.
- If you want to preview the modifications, click **Preview**.

Step 8 After the parameters have been modified, click **Change History** to view parameter modification details. For details, see **Viewing Change History of DB Instance Parameters**.

NOTICE

Check the value in the **Effective upon Restart** column. If it is set to:

- **Yes**: If an instance status on the **Instances** page is **Pending restart**, the instance needs to be restarted to apply changes. If only one node in a replica set, shard, or config is restarted, the changes will not be applied.
- **No**: The changes are applied immediately.

----End

Modifying Parameters in a Custom Parameter Template

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Custom Templates**. Locate the target parameter template and click its name.

Step 6 Modify the required parameters.

For parameter description details, see Parameters.

- If you want to save the modifications, click **Save**.
- If you want to cancel the modifications, click **Cancel**.
- If you want to preview the modifications, click **Preview**.

Step 7 The modifications take effect only after you apply the parameter template to instances. For details, see **Applying a Parameter Template**.

NOTICE

- After the parameters have been modified, click **Change History** to view parameter modification details. For details, see **Viewing Change History of a Custom Parameter Template**.
- The change history page displays only the modifications of the last seven days.
- For details about the parameter template statuses, see **Parameter Template Status**.
- After modifying a parameter, view the associated instance status in the instance list. If **Pending restart** is displayed, restart the instance for the modification to take effect.

----End

10.4 Viewing Parameter Change History

You can view the change history of a parameter template.

Precautions

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

Viewing Change History of DB Instance Parameters

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the navigation pane on the left, choose **Parameters**. On the **Change History** tab, view the parameter name, original parameter value, new parameter value, modification status, and modification time.

----End

Viewing Change History of a Custom Parameter Template

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Parameter Templates** page, click **Custom Templates**. Locate the target parameter template and click its name.

Step 5 In the navigation pane on the left, choose **Change History**. Then, view the parameter name, original parameter value, new parameter value, modification status, and modification time.

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

----End

10.5 Exporting a Parameter Template

- You can export a parameter template of a DB instance for future use. You can also apply the exported parameter template to other instances by referring to **Applying a Parameter Template**.
- You can export the parameter template details (parameter names, values, and descriptions) of an instance to a CSV file for review and analysis.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Instances**. On the displayed page, click the target instance. The **Basic Information** page is displayed.

Step 5 In the navigation pane on the left, choose **Parameters**. On the **Parameters** tab, above the parameter list, click **Export**.

Figure 10-5 Exporting a parameter template

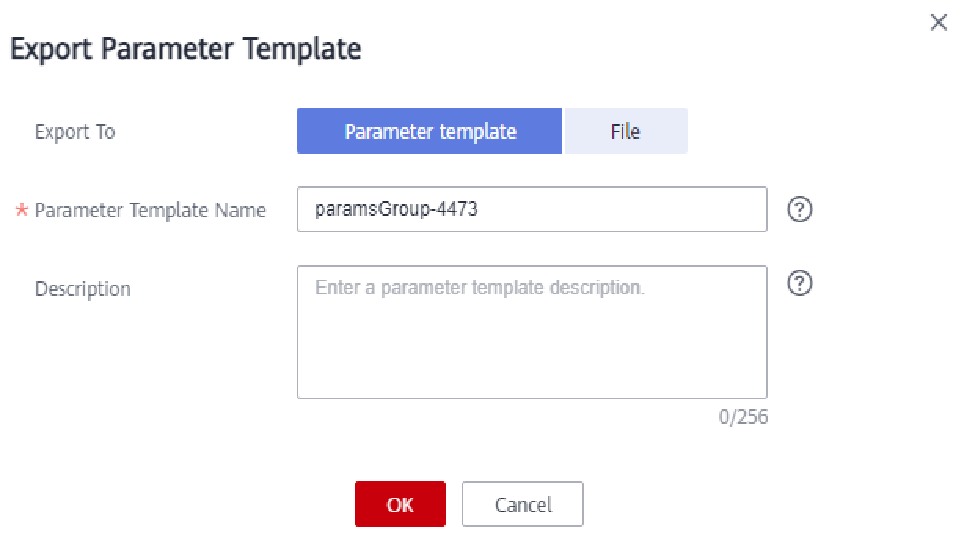

Export Parameter Template

Export To: Parameter template | File

* Parameter Template Name: paramsGroup-4473

Description: Enter a parameter template description.

0/256

- **Parameter Template**: The parameter list of the instance to will be exported to a parameter template for future use.

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

  NOTE

  - **New Parameter Template**: The template name can be up to 64 characters. It must start with a letter and can contain only letters (case-sensitive), digits, hyphens (-), periods (.), and underscores (_).
  - **Description**: It can contain up to 256 characters but cannot contain line breaks or the following special characters >!<"&'=

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

- **File**: The parameter template details (parameter names, values, and descriptions) of a DB instance are exported to a CSV file for review and analysis.

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

  NOTE

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

----End

10.6 Comparing Parameter Templates

This section describes how to compare two parameter templates of the same node type and DB engine version.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, locate the parameter template, and click **Compare**.

Step 6 In the displayed dialog box, select a parameter template that uses the same DB engine as the target template and click **OK**.

Figure 10-6 Comparing two parameter templates

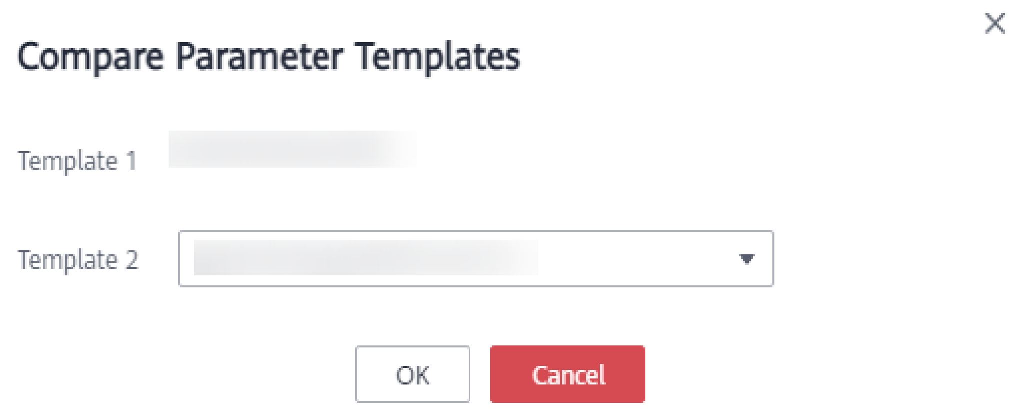

- If their settings are different, the parameter names and values of both parameter templates are displayed.
- If their settings are the same, no data is displayed.

----End

10.7 Replicating a Parameter Template

You can replicate a parameter template you have created. You can replicate a parameter template you have created. If you have a parameter template where you want to use most of its parameters and values in a new parameter template, you can create a replicate of template your existing template, or you can export a parameter template of a DB instance for future use.

Default parameter templates cannot be replicated, but you can create parameter templates based on them.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Custom Templates**, locate the parameter template, and click **Replicate** in the **Operation** column.

Step 6 In the displayed dialog box, enter the parameter template name and description and click **OK**.

Figure 10-7 Replicating a parameter template

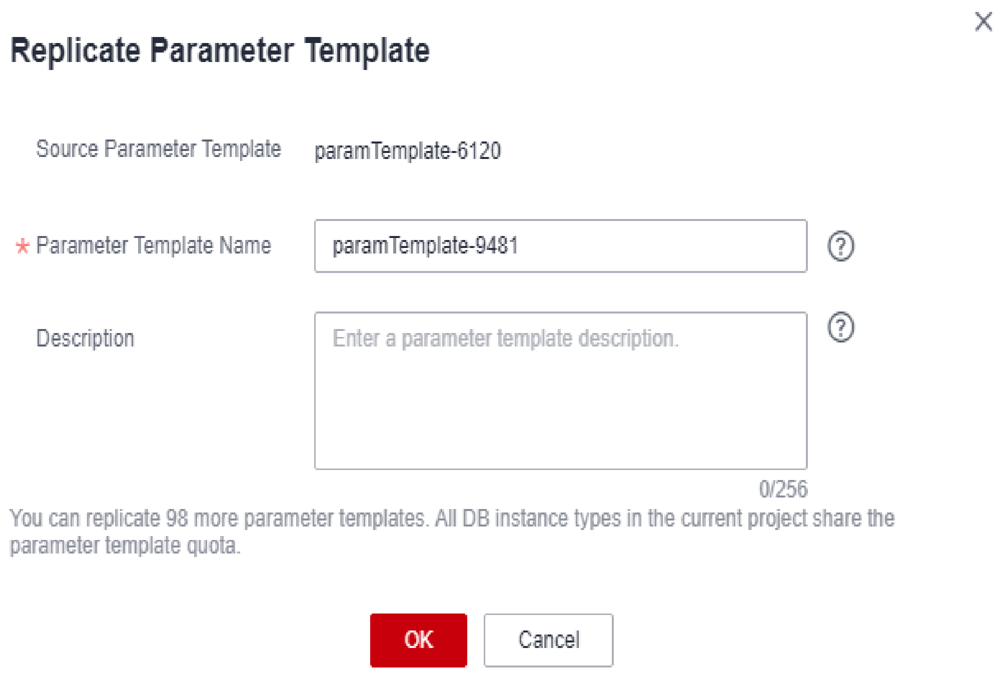

Replicate Parameter Template

Source Parameter Template paramTemplate-6120

* Parameter Template Name paramTemplate-9481

Description Enter a parameter template description.

0/256

You can replicate 98 more parameter templates. All DB instance types in the current project share the parameter template quota.

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

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

----End

10.8 Resetting a Parameter Template

This section describes how to reset all parameters in a parameter template you create to the default settings.

Precautions

Resetting a parameter template will restore all parameters in the parameter template to their default values. Exercise caution when performing this operation.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, click **Custom Templates**, locate the parameter template, and choose **More** > **Reset** in the **Operation** column.

Step 6 In the displayed dialog box, click **Yes**.

----End

10.9 Applying a Parameter Template

Modifications to parameters in a custom parameter template take effect for DB instances only after you have applied the template to the DB instances.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, apply a default template or a custom template to the DB instance:

- To apply a default template, click the **Default Templates** tab, locate the required parameter template, and click **Apply** in the **Operation** column.
- To apply a custom template, click **Custom Templates**, locate the parameter template, and in the **Operation** column, choose **More** > **Apply**.

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

Step 6 In the displayed dialog box, select the node or instance to which the parameter template will be applied and click **OK**.

After the parameter template is successfully applied, you can view the application records by referring to **Viewing Application Records of a Parameter Template**.

----End

10.10 Viewing Application Records of a Parameter Template

You can view the application records of a parameter template.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, select the parameter template for which you want to view application records.

- Click **Default Templates**. Locate the parameter template and click **View Application Record**.
- Click **Custom Templates**. Locate the parameter template and choose **More** > **View Application Record**.

Step 6 You can view the name or ID of the DB instance that the parameter template applies to, as well as the application status, application time, and the causes of any failures that have occurred.

----End

10.11 Modifying the Description of a Parameter Template

The section describes how to modify the description of a parameter template you created so that you can distinguish and identify parameter templates.

Precautions

The description of a default parameter template cannot be modified.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, locate the parameter template, and click ✎ in the **Description** column.

Step 6 Enter new description information. The parameter template description can contain up to 256 characters but cannot contain line breaks or the following special characters >!<"&'=

- To submit the change, click ✓. After the modification is successful, you can view the new description in the **Description** column of the parameter template list.
- To cancel the change, click ✕.

----End

10.12 Deleting a Parameter Template

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

Precautions

- Default parameter templates and parameter templates applied to instances cannot be deleted.
- Deleted parameter templates cannot be restored. Exercise caution when performing this operation.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Parameter Templates**.

Step 5 On the **Parameter Templates** page, locate the parameter template you want to delete, and choose **More** > **Delete**.

Step 6 In the displayed dialog box, click **Yes**.

----End

11 Connection Management

11.1 Querying DB Instance Connections and Managing Sessions

Scenario

- You can query the number of internal and external connections of a DB instance and the source IP addresses of these connections.
- You can also manage the sessions of an instance node and kill an abnormal session that takes a long time.

Precautions

- This function is not available to ECS-hosted instances and instances in the creating, frozen, or abnormal state.
- Exercise caution when killing a session. Your operations will be recorded in logs.
- This function is available for replica set instances and cluster instances of version 3.4 or later.
- When the CPU usage reaches the upper limit, requests to kill sessions may time out. In this case, you have to try more than once.

Querying the Number of Connections of a DB Instance

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Data Admin Service**.

Figure 11-1 Data Admin Service

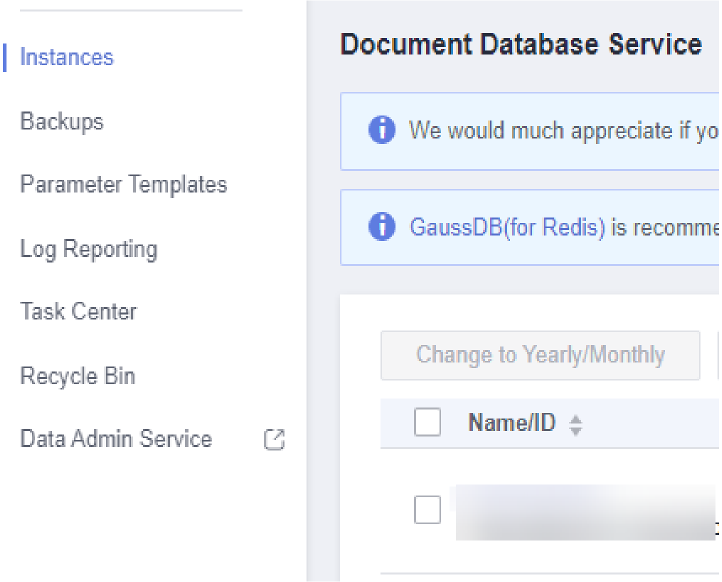

We would much appreciate if yo

GaussDB(for Redis) is recomme

Step 5 In the navigation pane on the left, choose **Intelligent O&M** > **Instance List**, and select **DDS DB Instances** from the drop-down list in the upper right corner of the page.

Figure 11-2 Instance List

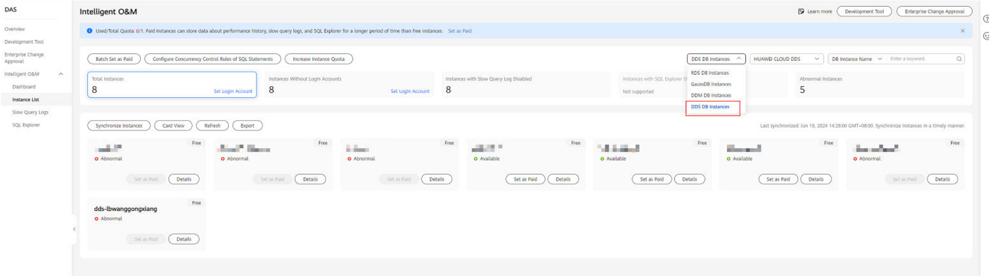

Used/Total Quota: 0/1. Paid instances can store data about performance history, slow query logs, and SQL Explorer for a longer period of time than free instances. Set as Paid

Last synchronized: Jun 19, 2024 14:28:00 GMT+08:00. Synchronize instances in a timely manner.

Step 6 Click **Details** in the area of the target instance.

Step 7 Click the **Connections** page to view the number of internal and external connections of the current DB instance and the source IP addresses of these connections.

Figure 11-3 Connections

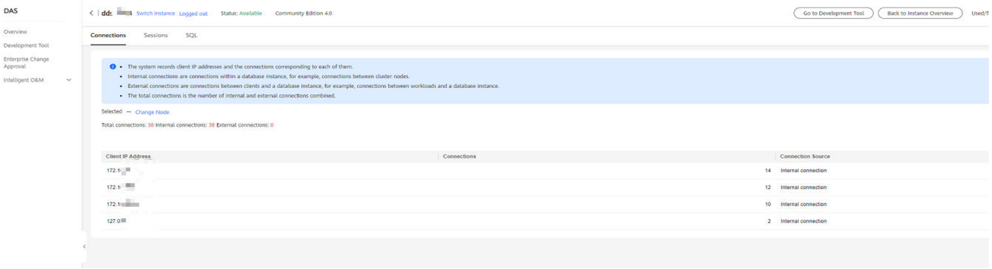

- The system records client IP addresses and the connections corresponding to each of them.
- Internal connections are connections within a database instance, for example, connections between cluster nodes.
- External connections are connections between clients and a database instance, for example, connections between workloads and a database instance.
- The total connections is the number of internal and external connections combined.

Total connections: 38 Internal connections: 38 External connections: 0

Client IP Address	Connections	Connection Source
172.1[redacted]	14	Internal connection
172.1[redacted]	12	Internal connection
172.1[redacted]	10	Internal connection
127.0[redacted]	2	Internal connection

Step 8 On the displayed page, click **Change Node** to view the number of internal and external connections of a specified node in the DB instance and the source IP addresses of these connections.

----End

Managing Sessions

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

Step 2 Click ⌖ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Data Admin Service**.

Figure 11-4 Data Admin Service

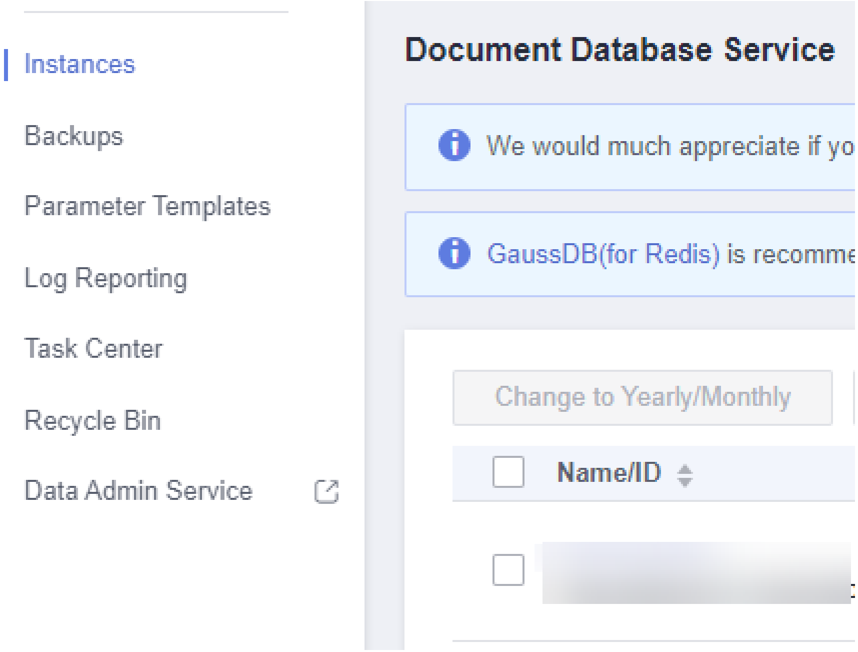

Document Database Service

We would much appreciate if yo

GaussDB(for Redis) is recomme

Step 5 In the navigation pane on the left, choose **Intelligent O&M** > **Instance List**, and select **DDS DB Instances** from the drop-down list in the upper right corner of the page.

Figure 11-5 Instance List

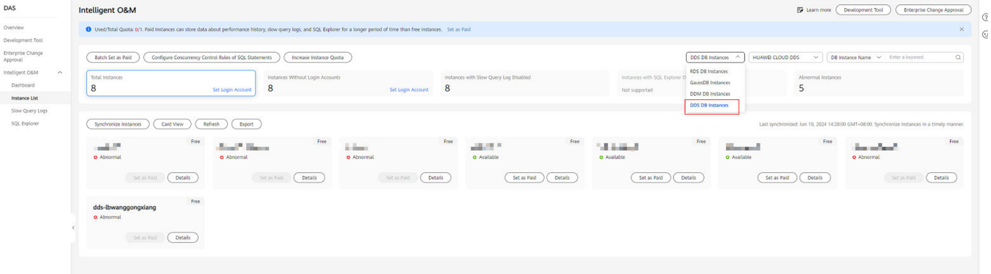

Used/Total Quota: 0/1. Paid instances can store data about performance history, slow query logs, and SQL Explorer for a longer period of time than free instances. Set as Paid

Total Instances 8 Set Login Account

Instances Without Login Accounts 8 Set Login Account

Instances with Slow Query Log Disabled 8

Instances with SQL Explorer D... Not supported

Abnormal Instances 5

Last synchronized: Jun 19, 2024 14:28:00 GMT+08:00. Synchronize instances in a timely manner.

Free Abnormal

Free Abnormal

Free Abnormal

Free Available

Free Available

Free Available

Free Abnormal

dds-lbwanggongxiang Free Abnormal

Step 6 Click the **Sessions** page to view the sessions of the current instance node.

Figure 11-6 Sessions

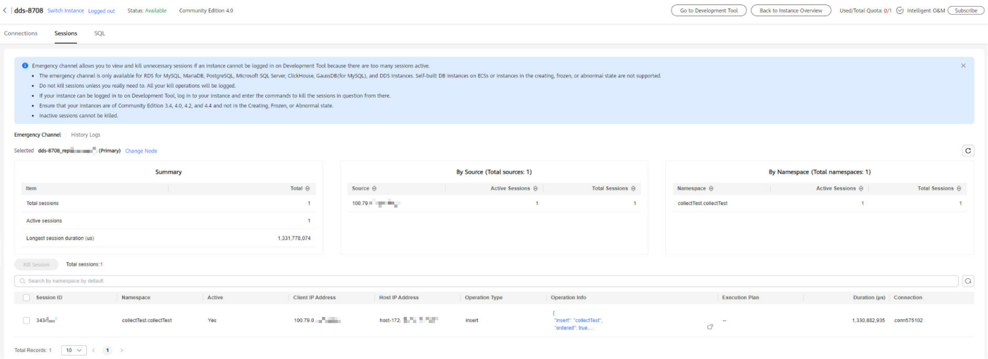

Emergency channel allows you to view and kill unnecessary sessions if an instance cannot be logged in on Development Tool because there are too many sessions active.

- The emergency channel is only available for RDS for MySQL, MariaDB, PostgreSQL, Microsoft SQL Server, ClickHouse, GaussDB(for MySQL), and DDS instances. Self-built DB instances on ECSs or instances in the creating, frozen, or abnormal state are not supported.
- Do not kill sessions unless you really need to. All your kill operations will be logged.
- If your instance can be logged in to on Development Tool, log in to your instance and enter the commands to kill the sessions in question from there.
- Ensure that your instances are of Community Edition 3.4, 4.0, 4.2, and 4.4 and not in the Creating, Frozen, or Abnormal state.
- Inactive sessions cannot be killed.

Selected dds-8708_rep[redacted] (Primary)

Summary

Item	Total
Total sessions	1
Active sessions	1
Longest session duration (us)	1,331,778,074

By Source (Total sources: 1)

Source	Active Sessions	Total Sessions
100.79[redacted]	1	1

By Namespace (Total namespaces: 1)

Namespace	Active Sessions	Total Sessions
collectTest.collectTest	1	1

Total sessions: 1

Search by namespace by default.

☐	Session ID	Namespace	Active	Client IP Address	Host IP Address	Operation Type	Operation Info	Execution Plan	Duration (μs)	Connection
☐	3434[redacted]	collectTest.collectTest	Yes	100.79.0.[redacted]	host-172[redacted]	insert	{ "insert": "collectTest", "ordered": true,...	--	1,330,882,935	conn575102

Total Records: 1 10 1

In the statistical item list, you can view total sessions, running sessions, and longest running session time of the current instance node. You can also view total sessions and active sessions by host or namespace. In the session list, you can view session details and perform the following operations:

- Select the abnormal session you want to end and click **Kill** for your database to recover.

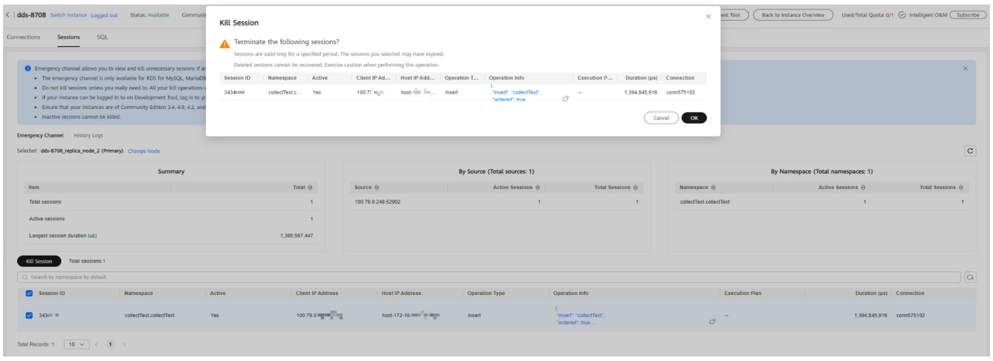

Kill Session

Terminate the following sessions?

Sessions are valid only for a specified period. The sessions you selected may have expired.

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

Session ID	Namespace	Active	Client IP Ad...	Host IP Add...	Operation T...	Operation Info	Execution P...	Duration (μs)	Connection
3434[redacted]	collectTest.c...	Yes	100.7[redacted]	host-[redacted]	insert	{ "insert": "collectTest", "ordered": true	--	1,394,845,919	conn575102

Emergency channel allows you to view and kill unnecessary sessions if ar
- The emergency channel is only available for RDS for MySQL, MariaDB
- Do not kill sessions unless you really need to. All your kill operations v
- If your instance can be logged in to on Development Tool, log in to yo
- Ensure that your instances are of Community Edition 3.4, 4.0, 4.2, and
- Inactive sessions cannot be killed.

Selected dds-8708_replica_node_2 (Primary)

Summary

Item	Total
Total sessions	1
Active sessions	1
Longest session duration (us)	1,395,567,447

By Source (Total sources: 1)

Source	Active Sessions	Total Sessions
100.79.0.248:52902	1	1

By Namespace (Total namespaces: 1)

Namespace	Active Sessions	Total Sessions
collectTest.collectTest	1	1

Total sessions: 1

Search by namespace by default.

☑	Session ID	Namespace	Active	Client IP Address	Host IP Address	Operation Type	Operation Info	Execution Plan	Duration (μs)	Connection
☑	343[redacted]	collectTest.collectTest	Yes	100.79.0[redacted]	host-172-16-[redacted]	insert	{ "insert": "collectTest", "ordered": true,...	--	1,394,845,919	conn575102

Total Records: 1 10 1

- Exercise caution when performing the kill operation. After you kill a session, you can switch to the **History Logs** tab to view the details.

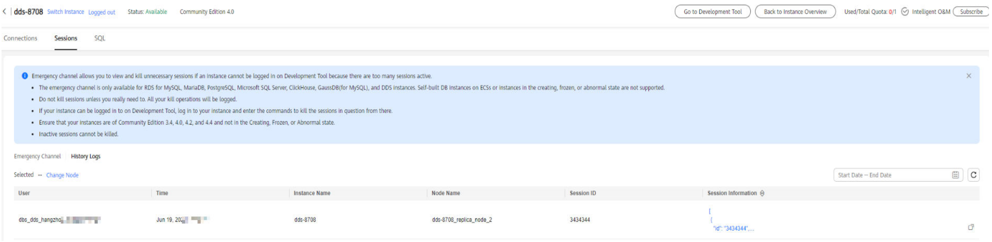

Emergency channel allows you to view and kill unnecessary sessions if an instance cannot be logged in on Development Tool because there are too many sessions active.

- The emergency channel is only available for RDS for MySQL, MariaDB, PostgreSQL, Microsoft SQL Server, ClickHouse, GaussDB(for MySQL), and DDS instances. Self-built DB instances on ECSs or instances in the creating, frozen, or abnormal state are not supported.
- Do not kill sessions unless you really need to. All your kill operations will be logged.
- If your instance can be logged in to on Development Tool, log in to your instance and enter the commands to kill the sessions in question from there.
- Ensure that your instances are of Community Edition 3.4, 4.0, 4.2, and 4.4 and not in the Creating, Frozen, or Abnormal state.
- Inactive sessions cannot be killed.

Selected --

User	Time	Instance Name	Node Name	Session ID	Session Information
dbs_dds_hangzho[redacted]	Jun 19, 20[redacted]	dds-8708	dds-8708_replica_node_2	3434344	[{ "id": "3434344",...

- Specify **Sessions lasting longer than** and click the refresh button to search for the sessions who last longer than this threshold.

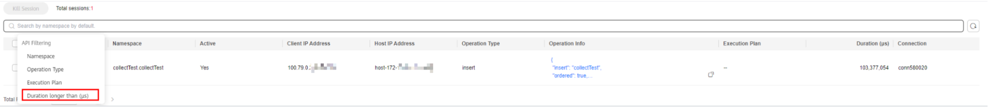

Total sessions: 1

Search by namespace by default.

☐	Namespace	Active	Client IP Address	Host IP Address	Operation Type	Operation Info	Execution Plan	Duration (μs)	Connection
☐	collectTest.collectTest	Yes	100.79.0.[redacted]	host-172-[redacted]	insert	{ "insert": "collectTest", "ordered": true,...	--	103,377,054	conn580020

----End

11.2 Configuring Cross-CIDR Access

When a replica set instance is connected through an internal network, a replica set node is configured with a management NIC (for receiving management instructions and internal communications of the instance) and a data NIC (for receiving and responding to service requests from the client), and the mapping between management IP addresses and data IP addresses of three standard CIDR blocks is configured by default.

- If your client and the replica set instance are deployed in different CIDR blocks and the client CIDR block is 192.168.0.0/16, 172.16.0.0/12, or 10.0.0.0/8, you do not need to configure **Access Across CIDR Blocks** for the instance.
- If your client and the replica set instance are deployed in different CIDR blocks and the client CIDR block is not 192.168.0.0/16, 172.16.0.0/12, or 10.0.0.0/8, you can configure **Access Across CIDR Blocks** for the instance to communicate with your client.
- No standard network segment is configured for replica set instances created before September 2021. If the client and the replica set instance are deployed in different network segments, you need to configure access across CIDR blocks to enable network connectivity.

This section describes how to configure **Access Across CIDR Blocks** for an instance.

Precautions

- Only replica set instances support this function.
- During the configuration of cross-CIDR access, services are running properly without interruption or intermittent disconnection.
- If the client and the replica set instance are in different VPCs and CIDR blocks, create a **VPC peering connection** between the VPCs and then configure cross-CIDR access.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Connections**.

Step 6 On the **Private Connection** tab, click **Enable** to the right of **Cross-CIDR Access**. You can add or delete the blocks as required.

- Click ⊕ to add new CIDR blocks.
- Click ⊖ to delete existing CIDR blocks.

Figure 11-7 Cross-CIDR Access

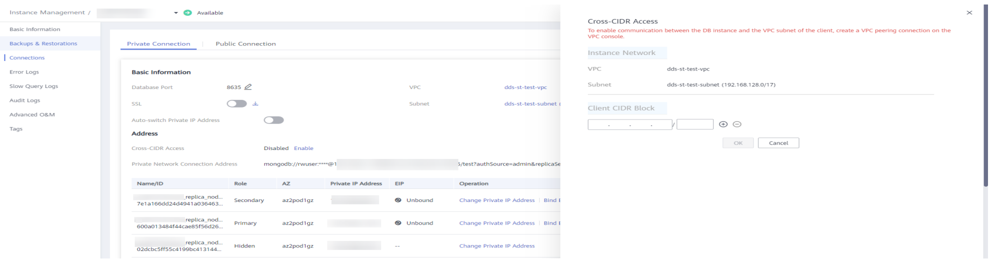

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet (

Auto-switch Private IP Address

Address

Cross-CIDR Access Disabled

Private Network Connection Address mongodb://rwuser:****@1[redacted]5/test?authSource=admin&replicaSe

Name/ID	Role	AZ	Private IP Address	EIP	Operation
[redacted]_replica_nod... 7e1a166dd24d4941a036463...	Secondary	az2pod1gz	[redacted]	Unbound	Change Private IP Address | Bind E
[redacted]replica_nod... 600a013484f44cae85f56d26...	Primary	az2pod1gz	[redacted]	Unbound	Change Private IP Address | Bind E
[redacted]replica_nod... 02dcbc5ff55c4199bc413144...	Hidden	az2pod1gz	[redacted]	--	Change Private IP Address

Cross-CIDR Access

To enable communication between the DB instance and the VPC subnet of the client, create a VPC peering connection on the VPC console.

Instance Network

VPC dds-st-test-vpc

Subnet dds-st-test-subnet (192.168.128.0/17)

Client CIDR Block

. . . /

NOTE

Up to 30 CIDR blocks can be configured, and each of them can overlap but they cannot be the same. That is, the source CIDR blocks can overlap but cannot be the same. The CIDR blocks cannot start with 127. The allowed IP mask ranges from 8 to 32.

Step 7 View the change results. After cross-CIDR access is enabled, **Enabled** is displayed to the right of **Cross-CIDR Access**.

If you need to change the client CIDR block, click **Change** to the right of **Cross-CIDR Access**.

Figure 11-8 Changing a CIDR block

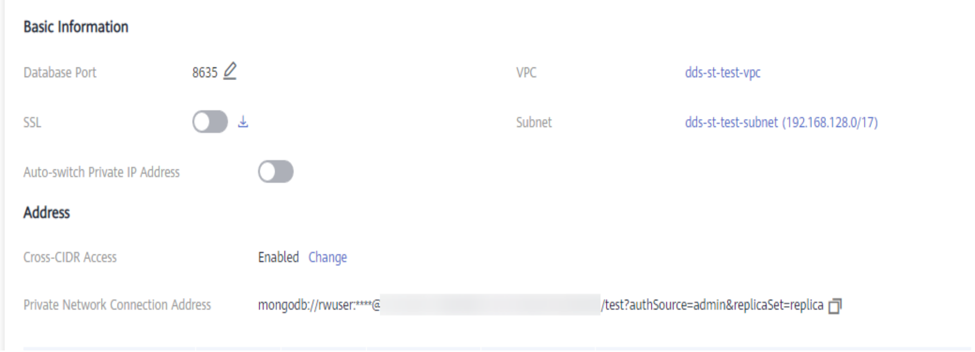

Basic Information

Database Port 8635

VPC dds-st-test-vpc

SSL

Subnet dds-st-test-subnet (192.168.128.0/17)

Auto-switch Private IP Address

Address

Cross-CIDR Access Enabled

Private Network Connection Address mongodb://rwuser:****@[redacted]/test?authSource=admin&replicaSet=replica

----End

Follow-up Operations

After cross-CIDR access is configured, you can use MongoShell to connect to a replica set instance over a private network. For details, see **Connecting to a Cluster Instance Using Mongo Shell**.

11.3 Enabling IP Addresses of Shard and Config Nodes

A cluster instance for Community Edition consists of dds mongos, shard, and config nodes. When your services need to read and write data from and into databases, connect to the dds mongos node. In certain scenarios (for example, data migration and synchronization between clusters), you need to read data from the shard or config node and will need to obtain the IP address of the corresponding node.

This section describes how to obtain the IP addresses of the shard and config nodes.

Before You Start

- If you need to use this function, choose **Service Tickets > Create Service Ticket** in the upper right corner of the management console and submit the application.
- DDS supports cluster instances of Community Edition 3.4, 4.0 and 4.2.
- DDS creates two connection addresses for the primary and secondary nodes in a shard group or config group.
- The network type of the connection address is the same as that of the current dds mongos node.
- Once the connection addresses are assigned to your nodes, they cannot be changed or deleted.
- If IPv6 is enabled in a subnet, you cannot enable IP addresses of the shard and config nodes for DB instances created using the subnet.
- After the shard IP address is enabled, restart the corresponding shard node for the configuration to take effect.
- After you enable the connection address, you can **connect to an instance using Mongo Shell**.

Enabling shard IP Address

NOTE

- The button for showing shard IP address can only be enabled. It cannot be disabled or modified.
- Once the shard IP address is enabled, DDS automatically applies for connection addresses for all shard nodes in the current instance.
- After the shard IP address is enabled and new shard nodes are added, you need to manually locate a newly added shard node and choose **More** > **Show shard IP Address** in the **Operation** column to show the shard IP address.
- After the shard IP address is enabled, the database user **sharduser** is created. For details about how to reset the password, see **Resetting the Password of User sharduser**.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Node Information** area, click the **shard** tab.

Figure 11-9 shard nodes

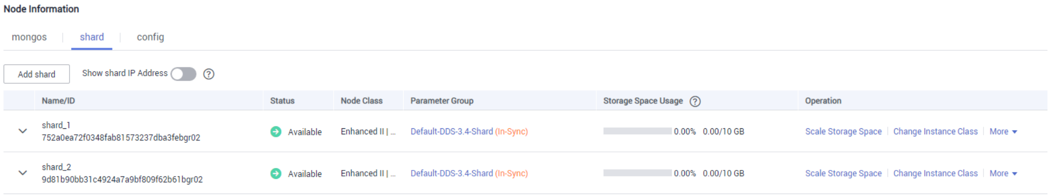

Node Information

Name/ID	Status	Node Class	Parameter Group	Storage Space Usage	Operation
shard_1 752a0ea72f0348fab81573237dba3febgr02	Available	Enhanced II | ...	Default-DDS-3.4-Shard (In-Sync)	0.00% 0.00/10 GB	Scale Storage Space | Change Instance Class | More
shard_2 9d81b90bb31c4924a7a9bf809f62b61bgr02	Available	Enhanced II | ...	Default-DDS-3.4-Shard (In-Sync)	0.00% 0.00/10 GB	Scale Storage Space | Change Instance Class | More

Step 6 Click **Show shard IP Address**. In the displayed dialog box, enter and confirm the password for connecting to the node.

Figure 11-10 Enable shard IP Address

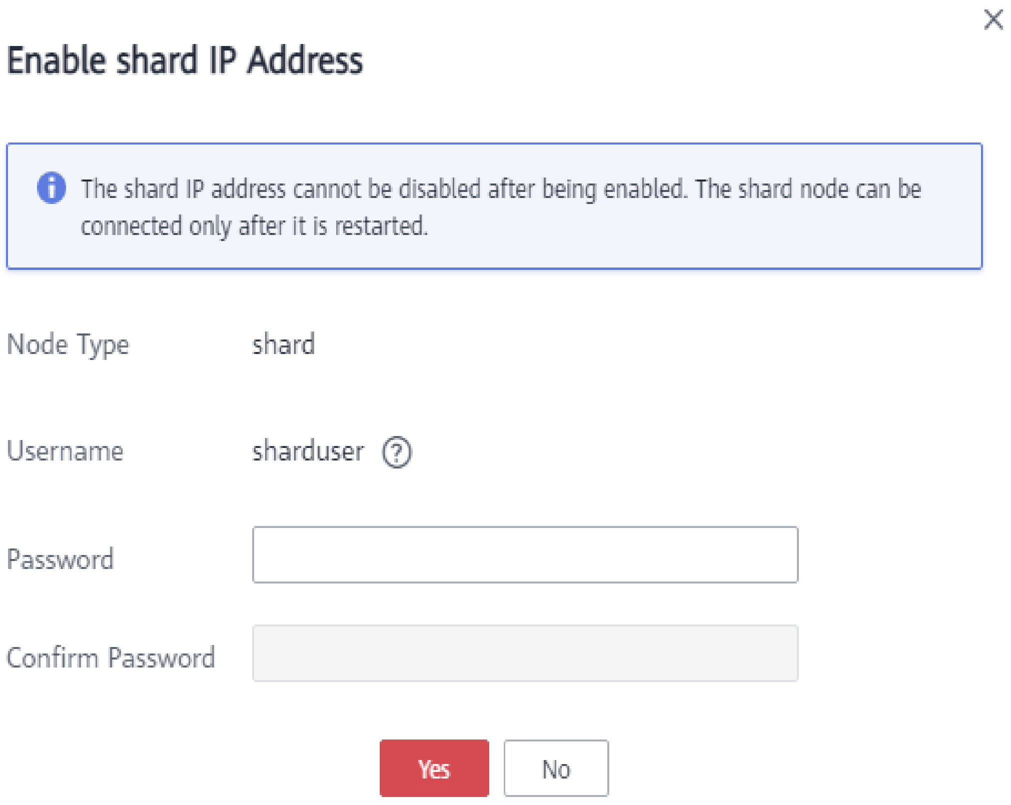

Enable shard IP Address

The shard IP address cannot be disabled after being enabled. The shard node can be connected only after it is restarted.

Node Type: shard

Username: sharduser

Password:

Confirm Password:

After the shard IP address is enabled, restart the corresponding shard node for the configuration to take effect.

In the **Node Information** area, locate the row that contains the shard node and click **Restart** in the **Operation** column to restart the shard node.

Figure 11-11 Restarting a shard node

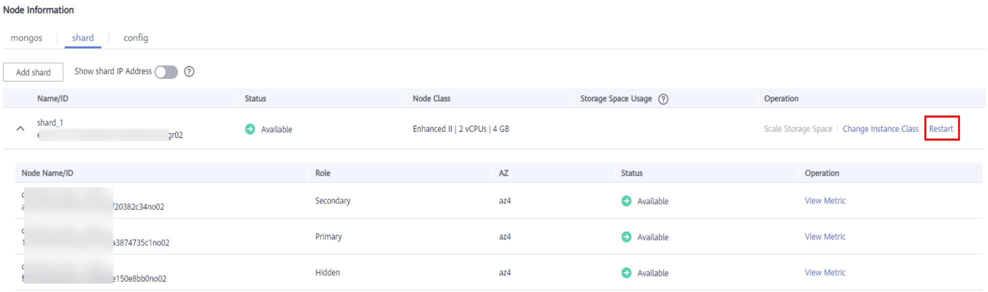

Node Information

mongos | shard | config

Add shard Show shard IP Address

	Name/ID	Status	Node Class	Storage Space Usage	Operation
^	shard_1 [redacted]gr02	Available	Enhanced II | 2 vCPUs | 4 GB		Scale Storage Space | Change Instance Class | Restart

Node Name/ID	Role	AZ	Status	Operation
[redacted]f20382c34no02	Secondary	az4	Available	View Metric
[redacted]a3874735c1no02	Primary	az4	Available	View Metric
[redacted]e150e8bb0no02	Hidden	az4	Available	View Metric

Step 7 View the private IP address of the shard node.

After the shard IP address is enabled, you can click ⌄ next to a shard node on the current page to expand the node drop-down list or click **Connections** in the navigation pane on the left, and then obtain the private IP address.

Figure 11-12 Private IP addresses of shard nodes

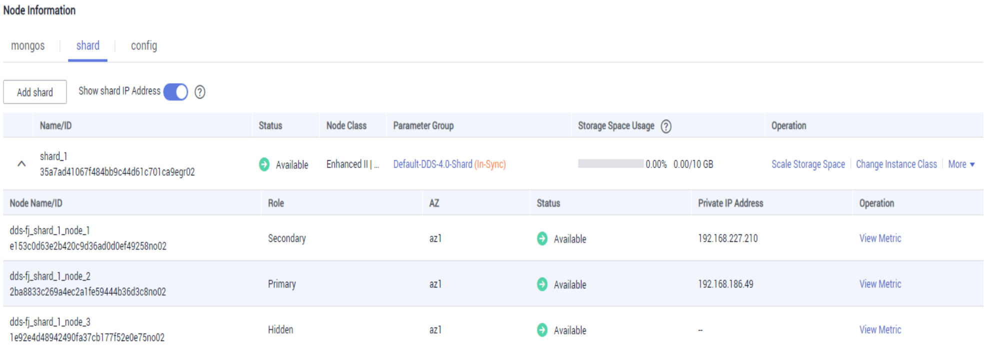

Node Information

mongos | shard | config

Add shard Show shard IP Address

	Name/ID	Status	Node Class	Parameter Group	Storage Space Usage	Operation
^	shard_1 35a7ad41067f484bb9c44d61c701ca9egr02	Available	Enhanced II | ...	Default-DDS-4.0-Shard (In-Sync)	0.00% 0.00/10 GB	Scale Storage Space | Change Instance Class | More ▾

Node Name/ID	Role	AZ	Status	Private IP Address	Operation
dds-fj_shard_1_node_1 e153c0d63e2b420c9d36ad0d0ef49258no02	Secondary	az1	Available	192.168.227.210	View Metric
dds-fj_shard_1_node_2 2ba8833c269a4ec2a1fe59444b36d3c8no02	Primary	az1	Available	192.168.186.49	View Metric
dds-fj_shard_1_node_3 1e92e4d48942490fa37cb177f52e0e75no02	Hidden	az1	Available	--	View Metric

The connection address of the current shard node is as follows:

mongodb://*sharduser:<password>@192.168.xx.xx:8637,192.168.xx.xx:8637*/test?authSource=admin&replicaSet=shard_?

NOTE

- **sharduser** is the username of the current shard node.
- ******** is the password of the current node.
- **192.168.xx.xx** and **192.168.xx.xx** are the private IP addresses of the primary and secondary shard nodes.
- **8637** is the port of the shard node and cannot be changed.
- **shard_?** is the name of the shard node to be connected, for example, **shard_1**.

----End

Resetting the Password of User sharduser

NOTE

This function is available only after the shard IP address is enabled.

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Node Information** area, click the **shard** tab.

Figure 11-13 shard nodes

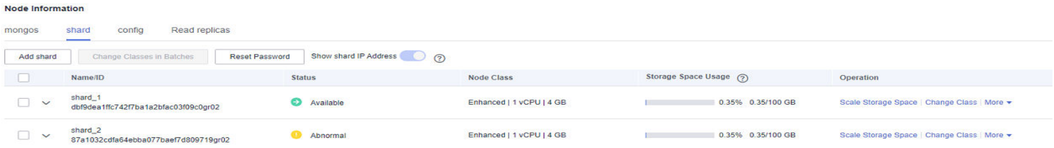

Node Information

	Name/ID	Status	Node Class	Storage Space Usage	Operation
□	shard_1 dbf9dea1ffc742f7ba1a2bfac03f09c0gr02	Available	Enhanced | 1 vCPU | 4 GB	0.35% 0.35/100 GB	Scale Storage Space | Change Class | More
□	shard_2 87a1032cdfa64ebba077baef7d809719gr02	Abnormal	Enhanced | 1 vCPU | 4 GB	0.35% 0.35/100 GB	Scale Storage Space | Change Class | More

Step 6 Click **Reset Password**.

Figure 11-14 Resetting a password

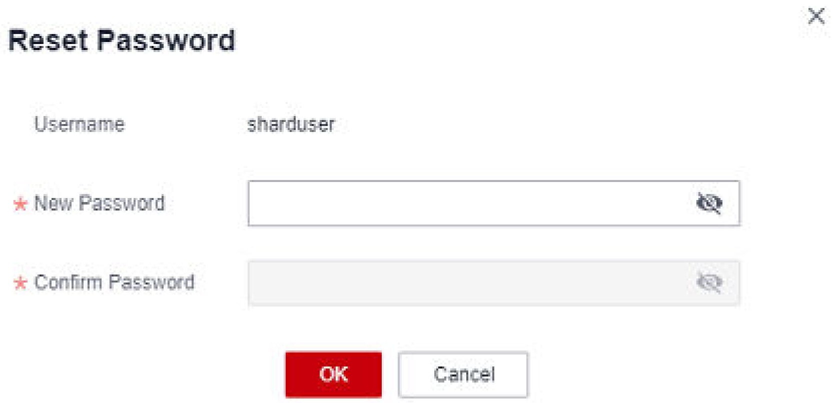

Reset Password

Username sharduser

* New Password

* Confirm Password

Step 7 Enter the new password and click **OK**.

----End

Enabling config IP Address

NOTE

- The button for showing config IP address can only be enabled. It cannot be disabled or modified.
- Once the config IP address is enabled, DDS automatically applies for connection addresses for all config nodes in the current instance.
- After the config IP address is enabled, the database user **csuser** is created. For details about how to reset the password, see **Resetting the Password of User csuser**.

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the left navigation pane, choose **Instances**. In the instance list, click the instance name to go to the **Basic Information** page.

Step 5 In the **Node Information** area, click the **config** tab.

Figure 11-15 config nodes

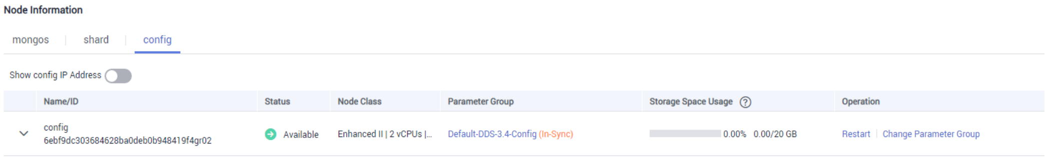

Node Information

Show config IP Address ○

	Name/ID	Status	Node Class	Parameter Group	Storage Space Usage ⓘ	Operation
⌄	config 6ebf9dc303684628ba0deb0b948419f4gr02	Available	Enhanced II | 2 vCPUs |...	Default-DDS-3.4-Config (In-Sync)	0.00% 0.00/20 GB	Restart | Change Parameter Group

Step 6 Click **Show config IP Address**. In the displayed dialog box, enter and confirm the password for connecting to the node.

Figure 11-16 Enable config IP Address

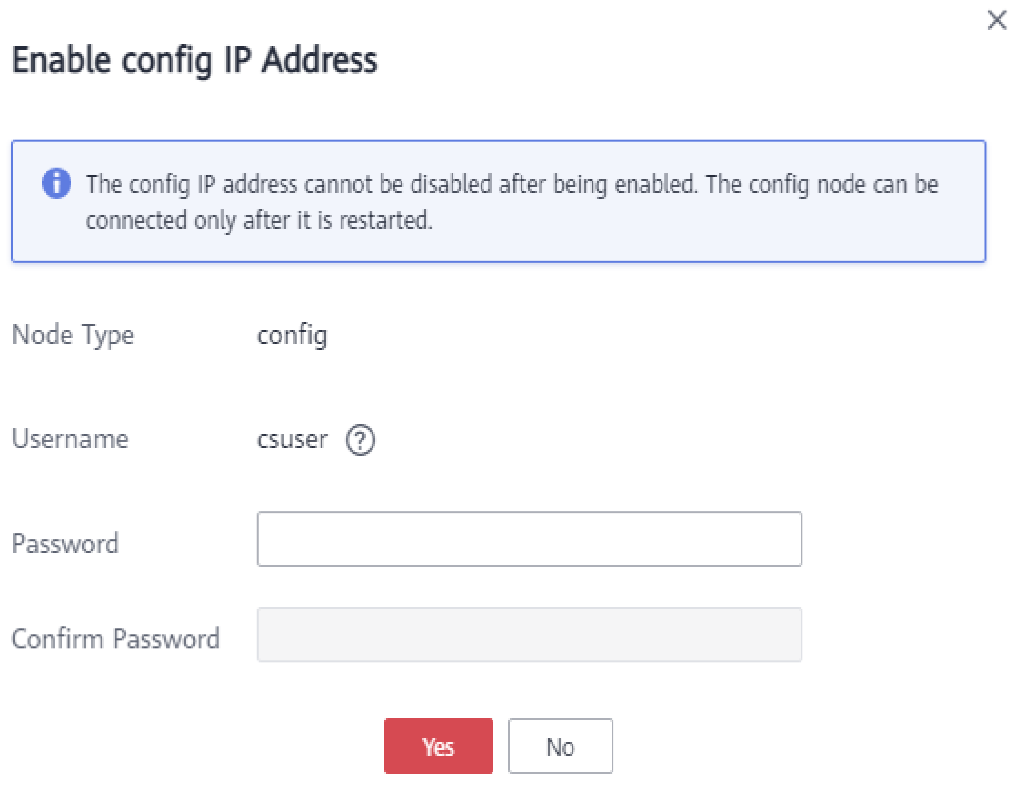

Enable config IP Address

The config IP address cannot be disabled after being enabled. The config node can be connected only after it is restarted.

Node Type config

Username csuser ⓘ

Password

Confirm Password

After the config IP address is enabled, the corresponding config node needs to be restarted for the configuration to take effect.

In the **Node Information** area, locate the row that contains the config node and click **Restart** in the **Operation** column to restart the config node.

Figure 11-17 Restarting a config node

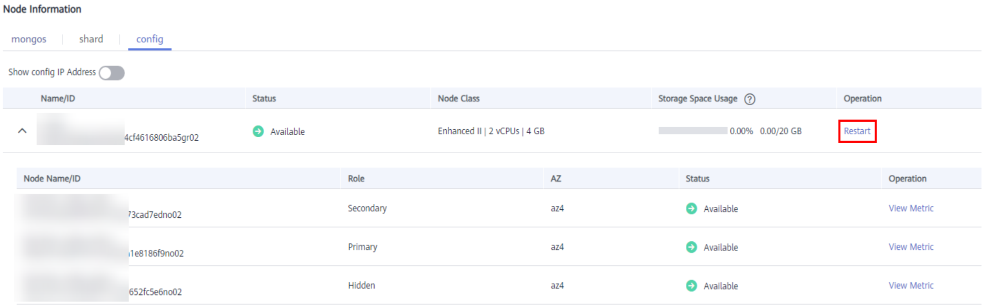

Node Information

	Name/ID	Status	Node Class	Storage Space Usage	Operation
^	[redacted] 4cf4616806ba5gr02	Available	Enhanced II | 2 vCPUs | 4 GB	0.00% 0.00/20 GB	Restart

Node Name/ID	Role	AZ	Status	Operation
[redacted] 73cad7edno02	Secondary	az4	Available	View Metric
[redacted] 1e8186f9no02	Primary	az4	Available	View Metric
[redacted] 652fc5e6no02	Hidden	az4	Available	View Metric

Step 7 View the private IP address of the config node.

After the config IP address is enabled, you can click ⌄ next to the node on the current page to expand the node drop-down list or click **Connections** in the navigation pane on the left, and then obtain the private IP address.

Figure 11-18 Private IP addresses of config nodes

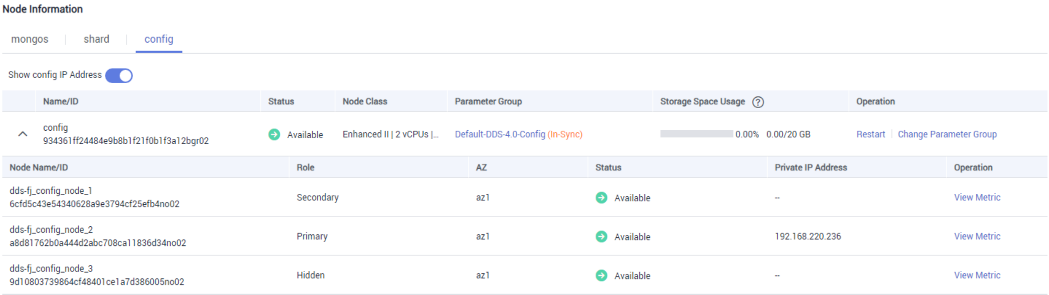

Node Information

	Name/ID	Status	Node Class	Parameter Group	Storage Space Usage	Operation
^	config 934361ff24484e9b8b1f21f0b1f3a12bgr02	Available	Enhanced II | 2 vCPUs |...	Default-DDS-4.0-Config (In-Sync)	0.00% 0.00/20 GB	Restart | Change Parameter Group

Node Name/ID	Role	AZ	Status	Private IP Address	Operation
dds-fj_config_node_1 6cfd5c43e54340628a9e3794cf25efb4no02	Secondary	az1	Available	--	View Metric
dds-fj_config_node_2 a8d81762b0a444d2abc708ca11836d34no02	Primary	az1	Available	192.168.220.236	View Metric
dds-fj_config_node_3 9d10803739864cf48401ce1a7d386005no02	Hidden	az1	Available	--	View Metric

The connection address of the current config node is as follows:

mongodb://***csuser:***<*password*>*@192.168.xx.xx:8636*/test?authSource=admin

NOTE

- **csuser** is the username of the current config node.
- ******** is the password of the current node.
- **192.168.xx.xx** is the private IP address of the primary config node.
- **8636** is the port of the config node and cannot be changed.

----End

Resetting the Password of User csuser

NOTE

This function is available only after the config IP address is enabled.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Step 5 In the **Node Information** area, click the **config** tab.

Figure 11-19 config nodes

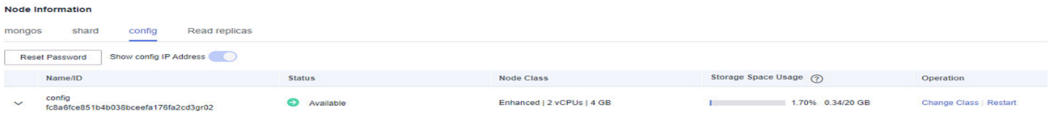

Node Information

mongos | shard | config | Read replicas

Reset Password | Show config IP Address

	Name/ID	Status	Node Class	Storage Space Usage	Operation
﹀	config fc8a6fce851b4b038bceefa176fa2cd3gr02	Available	Enhanced | 2 vCPUs | 4 GB	1.70% 0.34/20 GB	Change Class | Restart

Step 6 Click **Reset Password**.

Figure 11-20 Resetting a password

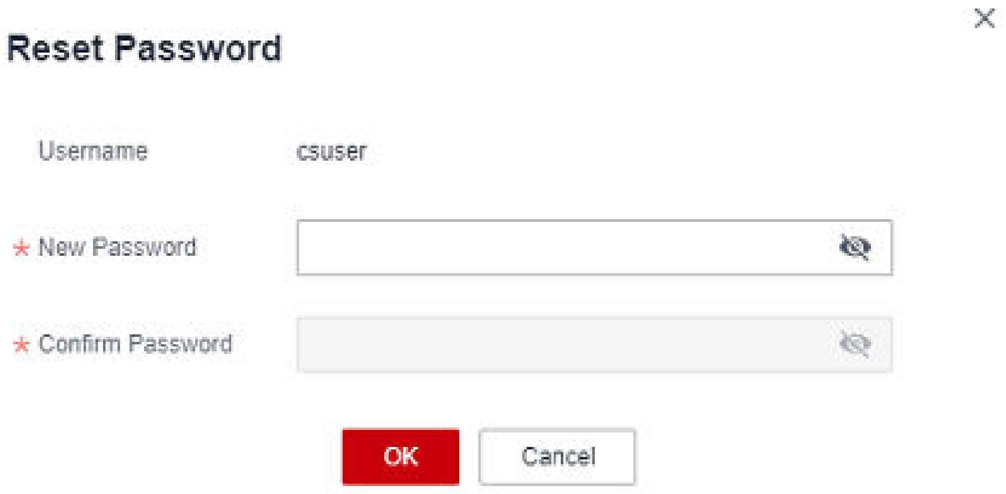

Reset Password

Username csuser

* New Password

* Confirm Password

Step 7 Enter the new password and click **OK**.

----End

Follow-up Operations

After the connection addresses of the shard or config nodes are enabled, you can connect to the shard or config nodes using MongoShell. The procedure is similar to that for connecting to a dds mongos node.For details, see **Connecting to a Cluster Instance Using Mongo Shell**.

11.4 Changing a Private IP Address

After data is migrated from an on-premises database or other cloud databases to DDS, the private IP address of the database may be changed. DDS allows you to change the private IP address, simplifying and accelerating the migration process.

Precautions

Changing the private IP address of a node will invalidate the previous private IP address. If an EIP is bound to the node, do not unbind the EIP during the change of the private IP address. After the change, the new private IP address is bound to the EIP.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name. The **Basic Information** page is displayed.

Alternatively, you can click **Connections** in the navigation pane on the left to go to the **Basic Information** page.

Step 5 In the **Node Information** area, locate the target node and click **Change Private IP Address** in the **Operation** column.

Step 6 In the displayed dialog box, enter a private IP address that is not in use and click **OK**.

Figure 11-21 Changing a private IP address

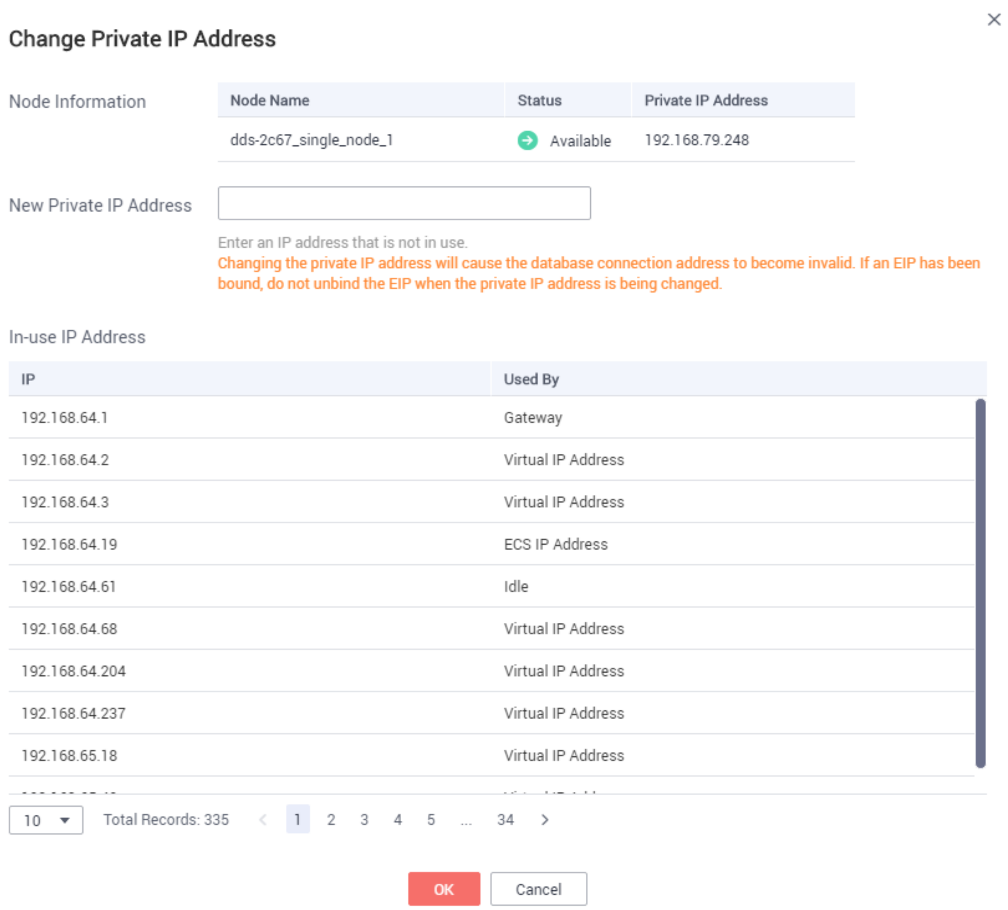

Change Private IP Address

Node Information

Node Name	Status	Private IP Address
dds-2c67_single_node_1	Available	192.168.79.248

New Private IP Address

Enter an IP address that is not in use.
Changing the private IP address will cause the database connection address to become invalid. If an EIP has been bound, do not unbind the EIP when the private IP address is being changed.

In-use IP Address

IP	Used By
192.168.64.1	Gateway
192.168.64.2	Virtual IP Address
192.168.64.3	Virtual IP Address
192.168.64.19	ECS IP Address
192.168.64.61	Idle
192.168.64.68	Virtual IP Address
192.168.64.204	Virtual IP Address
192.168.64.237	Virtual IP Address
192.168.65.18	Virtual IP Address

Step 7 In the **Node Information** area, locate the target node and view the new private IP address.

----End

11.5 Changing a Database Port

This section describes how to change a database port.

Precautions

- For security purposes, the database port cannot be modified when the instance is in any of the following statuses:
  - Frozen
  - Restarting
  - Adding node
  - Switching SSL
  - Changing instance class
  - Deleting node
  - The storage space is being expanded.
  - Abnormal
- The default port of a DB instance is 8635. After a DB instance is created, you can change its port number to a value ranging from 2100 to 65535 (excluding 12017 and 33071).

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the **Network Information** area on the **Basic Information** page, click ✎ in the **Database Port** field to change the database port.

Figure 11-22 Changing a database port

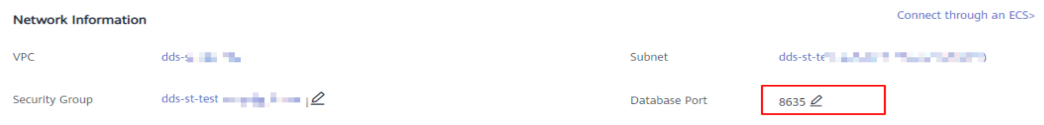

Network Information

Connect through an ECS>

VPC dds-s[redacted]

Subnet dds-st-te[redacted])

Security Group dds-st-test [redacted] | ✎

Database Port 8635 ✎

In the navigation pane on the left, choose **Connections** and click ✎ in the **Database Port** field in the **Basic Information** area to change the database port.

Figure 11-23 Changing a database port

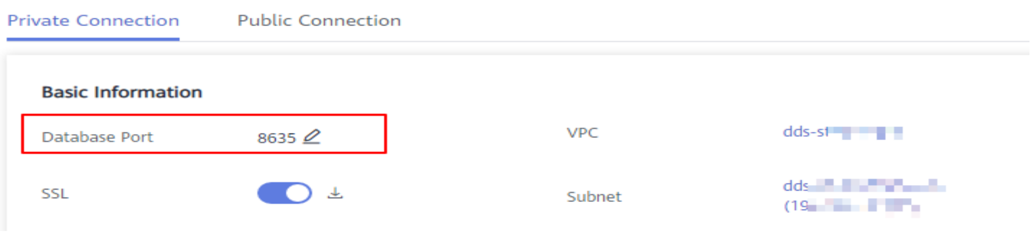

Basic Information

Database Port 8635

VPC dds-s[redacted]

SSL

Subnet dds[redacted]
(19[redacted]

📖 NOTE

The database port ranges from 2100 to 65535 (excluding 12017 and 33071).

- To submit the change, click ✓. This process takes about 1 to 5 minutes.
- To cancel the change, click ✗.

Step 6 View the modification result.

----End

11.6 Applying for and Modifying a Private Domain Name

You can apply for a private domain name and connect to DDS instances using the private domain name.

Precautions

- After a private domain name is generated, changing the private IP address will interrupt database connections. Exercise caution when performing this operation.
- You need to apply for the permissions needed to use private domain names. For details, contact customer service.
- When this function is enabled, you need to apply for a domain name for an existing instance. A domain name is automatically applied for a new instance.
- This function is available in the following regions: CN North-Beijing4, CN East-Shanghai1, CN South-Guangzhou, CN-Hong Kong, and CN Southwest-Guiyang1.

Applying for a Private Domain Name

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name and go to the **Basic Information** page.

Step 5 In the **Node Information** area on the **Basic Information** page, click **Private Domain Name**.

Alternatively, in the navigation pane on the left, choose **Connections**. In the **Basic Information** area on the **Private Connection** tab, click **Private Domain Name**.

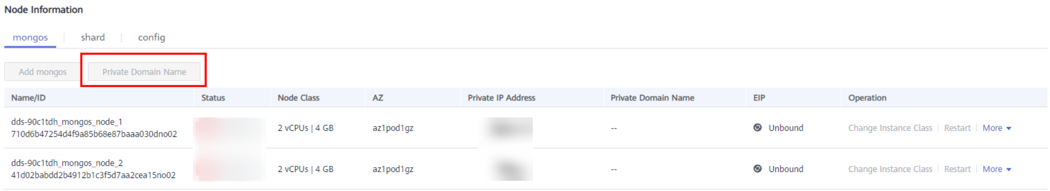

Node Information

mongos | shard | config

Add mongos | Private Domain Name

Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name	EIP	Operation
dds-90c1tdh_mongos_node_1 710d6b47254d4f9a85b68e87baaa030dno02	[redacted]	2 vCPUs | 4 GB	az1pod1gz	[redacted]	--	Unbound	Change Instance Class | Restart | More ▾
dds-90c1tdh_mongos_node_2 41d02babdd2b4912b1c3f5d7aa2cea15no02	[redacted]	2 vCPUs | 4 GB	az1pod1gz	[redacted]	--	Unbound	Change Instance Class | Restart | More ▾

Figure 11-24 Applying for a private domain name

Step 6 In the **Node Information** area on the **Basic Information** page, view the generated private domain names in the **Private Domain Name** column.

Alternatively, click **Connections** in the navigation pane on the left. In the **Basic Information** area on the displayed page, view the generated private domain names in the **Private Domain Name** column.

----End

Modifying a Private Domain Name

You can change the private domain name of an existing DB instance.

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the **Node Information** area on the **Basic Information** page, choose **More** > **Change Private Domain Name** in the **Operation** column.

Alternatively, choose **Connections** in the navigation pane on the left. In the lower part of the **Basic Information** area on the **Private Connection** tab, choose **More** > **Change Private Domain Name** in the **Operation** column.

Step 6 In the displayed dialog box, enter a new private domain name. Click **OK**. After the private domain name is changed, it takes about 5 minutes for the change to take effect.

NOTE

- Only the prefix of a private domain name can be modified.
- The prefix of a private domain name can contain 8 to 56 characters, and can include only letters and digits.
- The new private domain name must be different from the existing ones.

Step 7 If you have enabled the operation protection function, click **Send Code** in the displayed **Identity Verification** dialog box and enter the obtained verification code. Then, click **OK**.

Two-factor authentication improves the security of your account and cloud product. For details about how to enable operation protection, see the *Identity and Access Management User Guide*.

----End

12 Database Usage

12.1 Creating a Database Account Using Commands

When you create a DDS instance, the system automatically creates the default account **rwuser**. You can use the default account **rwuser** to create other database accounts based on service requirements. Then, you can use the default account **rwuser** or other created accounts to perform operations on data in the database, such as databases, tables, and indexes.

Precautions

- When creating a database account for a specified instance, you are advised to enable SSL to improve data security.
- If the existing DDS instances are of version 3.2, you cannot create database accounts for them. You can only change the password of the administrator account **rwuser**.
- When creating a database account, configure **passwordDigestor:"server"**. For details, see the **official document**.

Prerequisites

A DDS instance has been connected. For details, see "Connecting to an Instance over a Public Network" and "Connecting to an Instance over a Private Network" in *Document Database Service Getting Started.*

Account Description

- When a DDS instance is created, users **root**, **monitor**, and **backup** are automatically created. These accounts belong to the Huawei Cloud DB instance management platform and cannot be operated or used. Attempting to delete, rename, change the passwords, or change privileges for these accounts will result in errors.
- You can change the password of the database administrator **rwuser** and any accounts you create.

- The default user **rwuser** and users created by **rwuser** have limited permissions on system databases **admin** and **config**. They have all required permissions on the databases and tables created under them.
- Generally, a MongoDB user is created in a specified authentication database. When connecting to a database, use **--authenticationDatabase** to specify the corresponding authentication database.
- In a DDS instance, the default authentication database of user **rwuser** is **admin**.
- If you enter incorrect passwords for five consecutive times, the account will be locked for 10s.

Setting Password Strength for Database Accounts

- The administrator password must meet the following password policy:
  - Contains 8 to 32 characters.
  - Must be a combination of uppercase letters, lowercase letters, digits, and special characters: ~!@#%^*-_=+?()$
- The database user created on the client must meet the following password policy:
  - Contains 8 to 32 characters.
  - Must be a combination of uppercase letters, lowercase letters, digits, and special characters: ~@#%-_!*+=^?

When you create a DB instance or set a password, DDS automatically checks your password strength. If the password does not meet the complexity requirements, change the password as prompted.

Creating an Account

Step 1 Run the following command to select the admin database:

use admin

Step 2 Run the following command to create a database account (**user1** as an example):

db.createUser({user: "user1", pwd: "***", passwordDigestor:"server", roles: [{role: "root", db: "admin"}]})**

- ***server*** indicates the password encrypted on the server. It has a fixed value and does not need to be changed.
- *****:* indicates the example new password. The password must be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and special characters, such as ~@#%-_!*+=^?
- ***roles*** restricts the rights of the account. If an empty array is specified, the account does not have any permission.

Step 3 Check the result:

The account is successfully created if the following information is displayed:

```
Successfully added user: {
        "user" : "user1",
        "passwordDigestor" : "server",
        "roles" : [
          {
               "role" : "root",
               "db" : "admin"
          }
     ]
}
```

----End

Changing a Password

Step 1 Run the following command to select the admin database:

use admin

Step 2 Uses user **user1** as an example. Run the following command to change its password:

db.updateUser("user1", {passwordDigestor:"server",pwd:"newPasswd12#"})

- ***server*** indicates the password encrypted on the server. It has a fixed value and does not need to be changed.
- ***newPasswd12#***: indicates the example new password. The password must be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and special characters, such as ~@#%-_!*+=^?
- If the password contains any of the special characters @/%?# and is used in the MongoDB URL, escape the special characters in the URL and replace them with hexadecimal URL codes (ASCII codes).

Step 3 Check the setting result. The password is successfully changed if the following information is displayed:

- Cluster
  ```
  mongos>
  ```
- Replica set
  ```
  replica:PRIMARY>
  ```
- Single node
  ```
  replica:PRIMARY>
  ```

----End

Connecting to an Instance Using the Created Account

After a database account is created, it can be used to connect to the database. The operation details are as follows:

- **Connecting to a Cluster Instance Using Mongo Shell (Private Network)**
- **Connecting to a Cluster Instance Using Mongo Shell (Public Network)**
- **Connecting to a Replica Set Instance Using Mongo Shell (Private Network)**
- **Connecting to a Replica Set Instance Using Mongo Shell (Public Network)**
- **Connecting to a Single Node Instance Using Mongo Shell (Private Network)**
- **Connecting to a Single Node Instance Using Mongo Shell (Public Network)**

12.2 Creating a Database Using Commands

A database is a collection of tables, indexes, views, stored procedures, and operators. To make it easier to manage DDS DB instances, you can create a database by running commands on the newly-created DB instance. If the database does not exist, create the database and switch to the new database. If the database exists, directly switch to the database.

Prerequisites

A DDS instance has been connected. For details, see "Connecting to an Instance over a Public Network" and "Connecting to an Instance over a Private Network" in *Document Database Service Getting Started.*

Procedure

Step 1 Create a database.

use *dbname*

dbname: indicates the name of the database to be created.

Figure 12-1 Creating databases

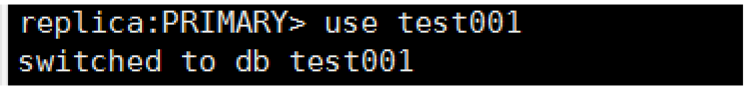

```
replica:PRIMARY> use test001
switched to db test001
```

Step 2 After a database is created, insert data into the database so that you can view the database in the database list.

Figure 12-2 Inserting data

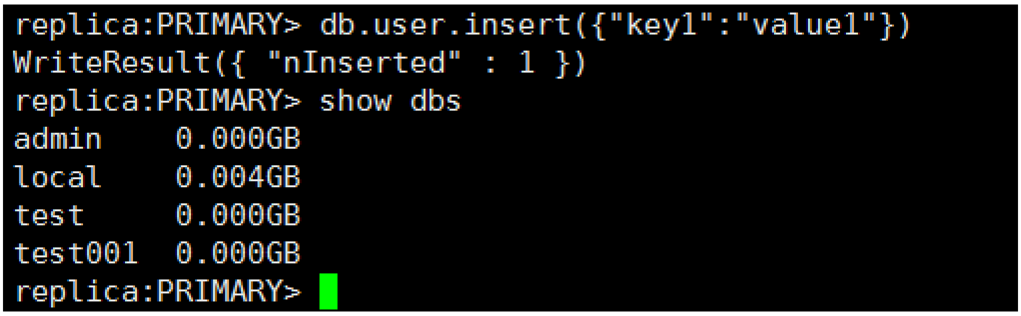

```
replica:PRIMARY> db.user.insert({"key1":"value1"})
WriteResult({ "nInserted" : 1 })
replica:PRIMARY> show dbs
admin    0.000GB
local    0.004GB
test     0.000GB
test001  0.000GB
replica:PRIMARY>
```

NOTE

There are three system databases created by default: admin, local, and test. If you directly insert data without creating a database, the data is inserted to the test database by default.

Figure 12-3 Viewing the database

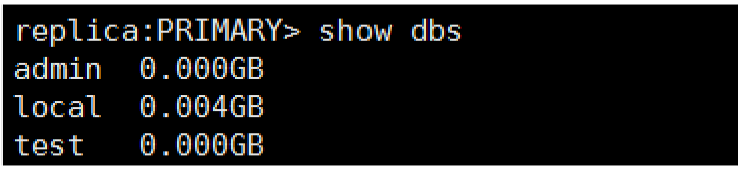

```
replica:PRIMARY> show dbs
admin  0.000GB
local  0.004GB
test   0.000GB
```

Step 3 View data in the database.

Figure 12-4 Viewing data

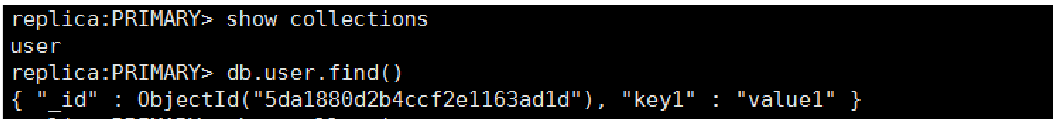

```
replica:PRIMARY> show collections
user
replica:PRIMARY> db.user.find()
{ "_id" : ObjectId("5da1880d2b4ccf2e1163ad1d"), "key1" : "value1" }
```

----End

12.3 Which Commands are Supported or Restricted by DDS?

The following tables list the commands supported and restricted by DDS.

For more information, see **official MongoDB documentation**.

NOTE

As shown in the following table, "√" indicates that the current version supports the command, and "x" indicates that the current version does not support the command.

Table 12-1 Commands supported and restricted by DDS

Type	Command	3.4	4.0	4.2	Description
Aggregates Commands	aggregate	√	√	√	-
	count	√	√	√	-
	distinct	√	√	√	-
	group	√	√	√	-
	mapReduce	√	√	√	This command can be used only when the **security.javascriptEnabled** parameter in the parameter template associated with the DB instance is set to **true**. For more information, see **How Do I Use MapReduce Commands?**

Type	Command	3.4	4.0	4.2	Description
Geospatial Commands	geoNear	√	√	√	-
	geoSearch	√	√	√	-
Query and Write Operation Commands	find	√	√	√	-
	insert	√	√	√	-
	update	√	√	√	-
	delete	√	√	√	-
	findAndModify	√	√	√	-
	getMore	√	√	√	-
	getLastError	√	√	√	-
	resetError	√	√	√	-
	getPrevError	√	√	√	-
	parallelCollectionScan	√	√	√	-
Query Plan Cache Commands	planCacheListFilters	√	√	√	-
	planCacheSetFilter	√	√	√	-
	planCacheClearFilters	√	√	√	-
	planCacheListQueryShapes	√	√	√	-
	planCacheListPlans	√	√	√	-
	planCacheClear	√	√	√	-
Authentication Commands	logout	√	√	√	-
	authenticate	√	√	√	-
	copydbgetnonce	√	√	√	-
	getnonce	√	√	√	-
	authSchemaUpgrade	x	x	x	System command

Type	Command	3.4	4.0	4.2	Description
User Management Commands	createUser	√	√	√	-
	updateUser	√	√	√	-
	dropUser	√	√	√	-
	dropAllUsersFromDatabase	√	√	√	-
	grantRolesToUser	√	√	√	-
	revokeRolesFromUser	√	√	√	-
	usersInfo	√	√	√	-
Role Management Commands	invalidateUserCache	√	√	√	-
	createRole	√	√	√	-
	updateRole	√	√	√	-
	dropRole	√	√	√	-
	dropAllRolesFromDatabase	√	√	√	-
	grantPrivilegesToRole	√	√	√	-
	revokePrivilegesFromRole	√	√	√	-
	grantRolesToRole	√	√	√	-
	revokeRolesFromRole	√	√	√	-
	rolesInfo	√	√	√	-
Replication Commands	replSetElect	x	x	x	System command
	replSetUpdatePosition	x	x	x	System command
	appendOplogNote	x	x	x	System command
	replSetFreeze	x	x	x	System command
	replSetGetStatus	√	√	√	-

Type	Command	3.4	4.0	4.2	Description
	replSetInitiate	x	x	x	System command
	replSetMaintenance	x	x	x	System command
	replSetReconfig	x	x	x	System command
	replSetStepDown	x	x	x	System command
	replSetSyncFrom	x	x	x	System command
	replSetRequestVotes	x	x	x	System command
	replSetDeclareElectionWinner	x	x	x	System command
	resync	x	x	x	System command
	applyOps	x	x	x	System command
	isMaster	√	√	√	-
	replSetGetConfig	x	x	x	System command
Sharding Commands	flushRouterConfig	√	√	√	High-risk commands
	addShard	x	x	x	Unauthorized operation
	addShardToZone	√	√	√	-
	balancerStart	√	√	√	-
	balancerStatus	√	√	√	-
	balancerStop	√	√	√	-
	removeShardFromZone	√	√	√	-
	updateZoneKeyRange	√	√	√	-
	cleanupOrphaned	x	x	x	High-risk commands

Type	Command	3.4	4.0	4.2	Description
	checkShardingIndex	x	x	x	System command
	enableSharding	√	√	√	-
	listShards	x	x	x	System command
	removeShard	x	x	x	High-risk commands
	getShardMap	x	x	x	System command
	getShardVersion	√	√	√	-
	mergeChunks	√	√	√	-
	setShardVersion	x	x	x	System command
	shardCollection	√	√	√	-
	shardingState	x	x	x	System command
	unsetSharding	x	x	x	System command
	split	√	√	√	-
	splitChunk	√	√	√	-
	splitVector	√	√	√	-
	moveChunk	√	√	√	-
	movePrimary	√	x	√	-
	isdbgrid	√	√	√	-
Administration Commands	setFeatureCompatibilityVersion	√	√	√	-
	renameCollection	√	√	√	-
	dropDatabase	√	√	√	-
	listCollections	√	√	√	-
	drop	√	√	√	-
	create	√	√	√	-

Type	Command	3.4	4.0	4.2	Description
	clone	x	x	x	System command
	cloneCollection	√	√	√	-
	cloneCollectionAsCapped	√	√	√	-
	convertToCapped	√	√	√	-
	filemd5	√	√	√	-
	createIndexes	√	√	√	-
	listIndexes	√	√	√	-
	dropIndexes	√	√	√	-
	fsync	√	√	√	-
	clean	x	x	x	System command
	connPoolSync	x	x	x	System command
	connectionStatus	√	√	√	-
	compact	x	x	x	High-risk commands
	collMod	√	√	√	-
	reIndex	√	√	√	-
	setParameter	x	x	x	System configuration command
	getParameter	√	√	√	-
	repairDatabase	x	x	x	High-risk commands
	repairCursor	x	x	x	System command
	touch	√	√	√	-
	shutdown	x	x	x	High-risk commands
	logRotate	x	x	x	High-risk commands

Type	Command	3.4	4.0	4.2	Description
	killOp	√	√	√	-
	releaseFreeMemory	√	√	√	-
Diagnostic Commands	availableQueryOptions	√	√	√	-
	buildInfo	√	√	√	-
	collStats	√	√	√	-
	connPoolStats	x	x	x	System command
	cursorInfo	x	x	x	System command
	dataSize	√	√	√	-
	dbHash	x	x	x	System command
	dbStats	√	√	√	-
	diagLogging	x	x	x	System command
	driverOIDTest	x	x	x	System command
	explain	√	√	√	-
	features	√	√	√	-
	getCmdLineOpts	x	x	x	System command
	getLog	x	x	x	System command
	hostInfo	x	x	x	System command
	isSelf	x	x	x	System command
	listCommands	√	√	√	-
	listDatabases	√	√	√	-
	netstat	x	x	x	System command
	ping	√	√	√	-
	profile	√	√	√	-

Type	Command	3.4	4.0	4.2	Description
	serverStatus	√	√	√	-
	shardConnPoolStats	x	x	x	System command
	top	√	√	√	-
	validate	x	x	x	System configuration command
	whatsmyuri	√	√	√	-
Internal Commands	handshake	x	x	x	System command
	_recvChunkAbort	x	x	x	System command
	_recvChunkCommit	x	x	x	System command
	_recvChunkStart	x	x	x	System command
	_recvChunkStatus	x	x	x	System command
	_replSetFresh	x	x	x	System command
	mapreduce.shardedfinish	x	x	x	System command
	_transferMods	x	x	x	System command
	replSetHeartbeat	x	x	x	System command
	replSetGetRBID	x	x	x	System command
	_migrateClone	x	x	x	System command
	replSetElect	x	x	x	System command
	writeBacksQueued	x	x	x	System command
	writebacklisten	x	x	x	System command

Type	Command	3.4	4.0	4.2	Description
System Events Auditing Commands	logApplicationMessage	x	x	x	System command

13 Data Security

13.1 Enabling or Disabling SSL

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

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

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

Precautions

- Enabling or disabling SSL will cause instances to restart. Exercise caution when performing this operation.

  NOTE

  When you enable or disable SSL, DDS will restart once. During the restart, each node will be intermittently disconnected for about 30 seconds. You are advised to enable or disable SSL during off-peak hours and ensure that your applications support automatic reconnection.

- If SSL is enabled, you can connect to a database using SSL, which is more secure.

  Currently, insecure encryption algorithms are disabled. The following table lists the supported TLS versions and cipher suites.

Version	TLS Version	Cipher Suites
3.4	TLS 1.2	AES256-GCM-SHA384 AES128-GCM-SHA256
4.0	TLS 1.2	DHE-RSA-AES256-GCM-SHA384 DHE-RSA-AES128-GCM-SHA256

The server where the client is located must support the corresponding TLS version and encryption algorithm suite. Otherwise, the connection fails.

- If SSL is disabled, you can connect to a database using an unencrypted connection.

Enabling SSL

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the target DB instance.

Step 5 In the **DB Information** area on the **Basic Information** page, click next to the **SSL** field.

Figure 13-1 Enabling SSL

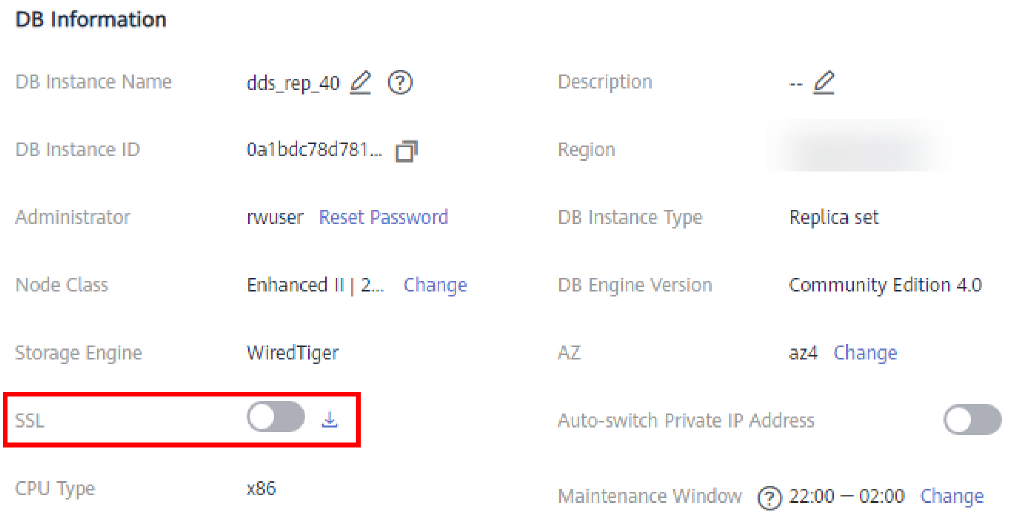

DB Information

DB Instance Name	dds_rep_40	Description	--
DB Instance ID	0a1bdc78d781...	Region	[redacted]
Administrator	rwuser Reset Password	DB Instance Type	Replica set
Node Class	Enhanced II | 2... Change	DB Engine Version	Community Edition 4.0
Storage Engine	WiredTiger	AZ	az4 Change
SSL	○	Auto-switch Private IP Address	○
CPU Type	x86	Maintenance Window	22:00 — 02:00 Change

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

Figure 13-2 Enabling SSL

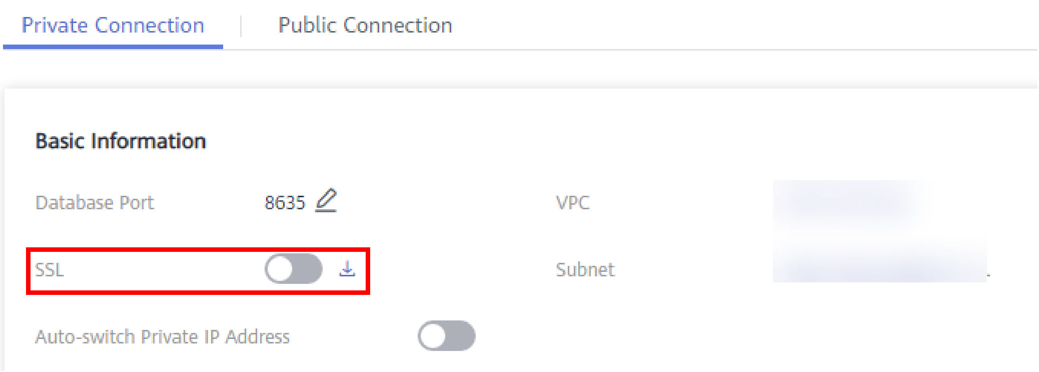

Private Connection | Public Connection

Basic Information

Database Port	8635	VPC	
SSL		Subnet	
Auto-switch Private IP Address			

Step 6 In the displayed dialog box, click **Yes**.

Step 7 In the **Basic Information** area, view the modification result.

Figure 13-3 SSL enabled

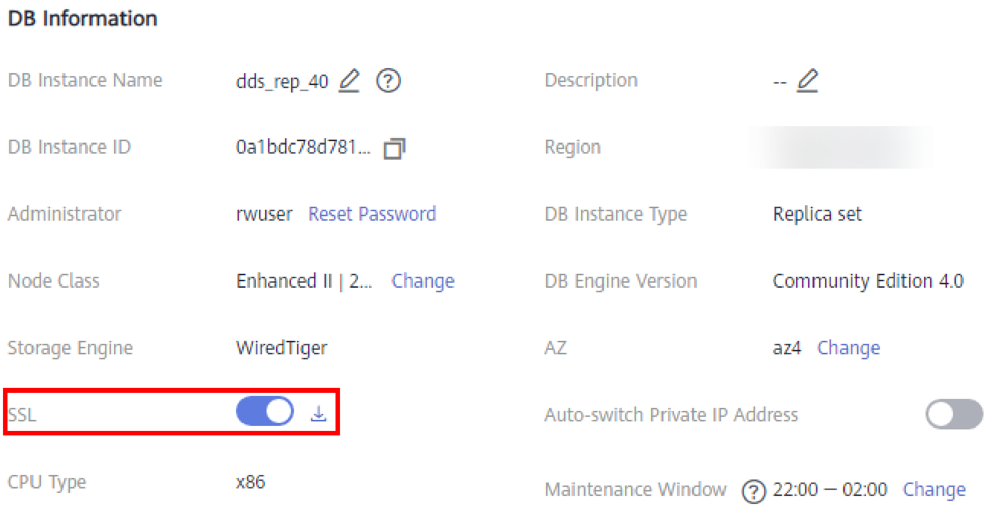

DB Information

DB Instance Name	dds_rep_40	Description	--
DB Instance ID	0a1bdc78d781...	Region	
Administrator	rwuser Reset Password	DB Instance Type	Replica set
Node Class	Enhanced II | 2... Change	DB Engine Version	Community Edition 4.0
Storage Engine	WiredTiger	AZ	az4 Change
SSL		Auto-switch Private IP Address	
CPU Type	x86	Maintenance Window	22:00 – 02:00 Change

Step 8 After SSL is enabled, click ⤓ next to **SSL** to download an SSL certificate.

For details about how to connect to an instance using SSL, refer to the following content:

- **Connecting to a Cluster Instance Using SSL**
- **Connecting to a Replica Set Instance Using SSL**
- **Connecting to a Single Node Instance Using SSL**

----End

Disabling SSL

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the target DB instance.

Step 5 In the **DB Information** area on the **Basic Information** page, click ● next to the **SSL** field.

Figure 13-4 Disabling SSL

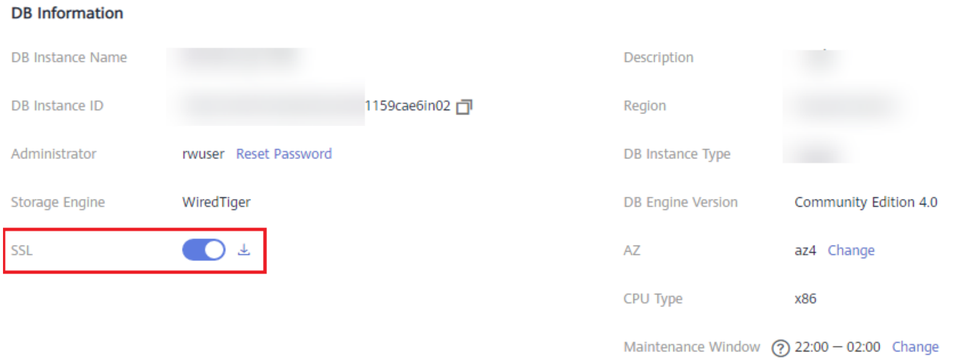

DB Information

DB Instance Name [redacted]
DB Instance ID [redacted]1159cae6in02
Administrator rwuser Reset Password
Storage Engine WiredTiger
SSL ●
Description [redacted]
Region [redacted]
DB Instance Type [redacted]
DB Engine Version Community Edition 4.0
AZ az4 Change
CPU Type x86
Maintenance Window 22:00 – 02:00 Change

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

Figure 13-5 Disabling SSL

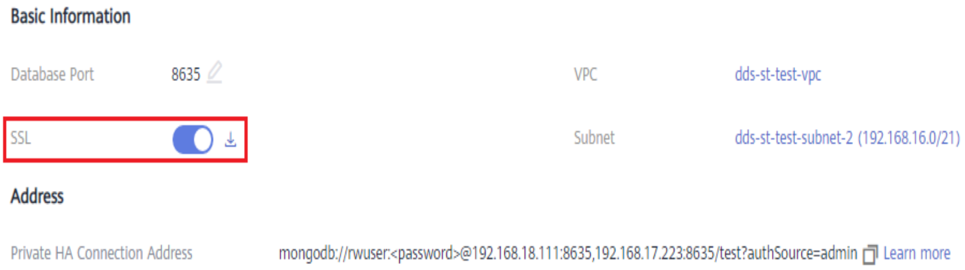

Basic Information

Database Port 8635
SSL ●
VPC dds-st-test-vpc
Subnet dds-st-test-subnet-2 (192.168.16.0/21)

Address

Private HA Connection Address mongodb://rwuser:<password>@192.168.18.111:8635,192.168.17.223:8635/test?authSource=admin Learn more

Step 6 In the displayed dialog box, click **Yes**.

Step 7 In the **Basic Information** area, view the modification result.

Figure 13-6 SSL disabled

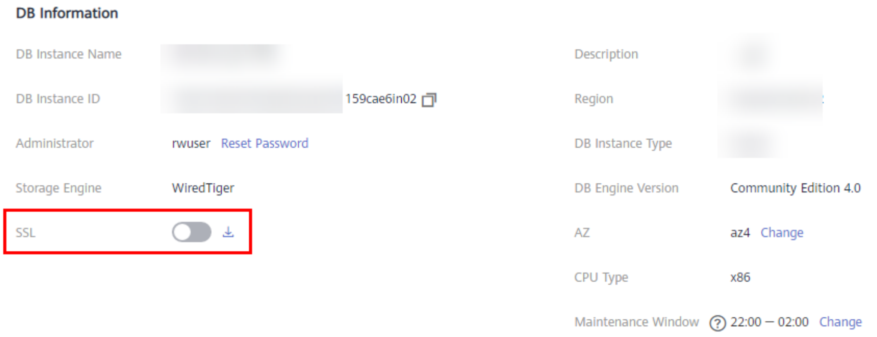

DB Information

DB Instance Name [redacted]
DB Instance ID [redacted]159cae6in02
Administrator rwuser Reset Password
Storage Engine WiredTiger
SSL ○
Description [redacted]
Region [redacted]
DB Instance Type [redacted]
DB Engine Version Community Edition 4.0
AZ az4 Change
CPU Type x86
Maintenance Window 22:00 – 02:00 Change

Step 8 Connect to an instance using an unencrypted connection.

For details, refer to the following content:

- **Connecting to a Cluster Instance Using an Unencrypted Connection**
- **Connecting to a Replica Set Instance Using an Unencrypted Connection**
- **Connecting to a Single Node Instance Using an Unencrypted Connection**

----End

13.2 Resetting the Administrator Password

For security reasons, you are advised to periodically change administrator passwords.

If you do not set the administrator password for the DB instance that you are creating, you need to reset the password before connecting to the DB instance.

Precautions

- You cannot reset the administrator password for an instance is in any of the following statuses:
  - Frozen
  - Creating
  - Restarting
  - Adding node
  - Switching SSL
  - Changing port
  - Changing instance class
  - Deleting node
  - Upgrading minor version
  - Switchover in progress
  - Changing AZ
  - Adding read replicas
- If you enable operation protection to improve the security of your account and cloud products, two-factor authentication is required for sensitive operations. For details about how to enable operation protection, see **Operation Protection** in *Identity and Access Management User Guide*.

⚠ CAUTION

Changing the password may interrupt services.

Procedure

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and choose **More** > **Reset Password** in the **Operation** column.

Figure 13-7 Resetting a password

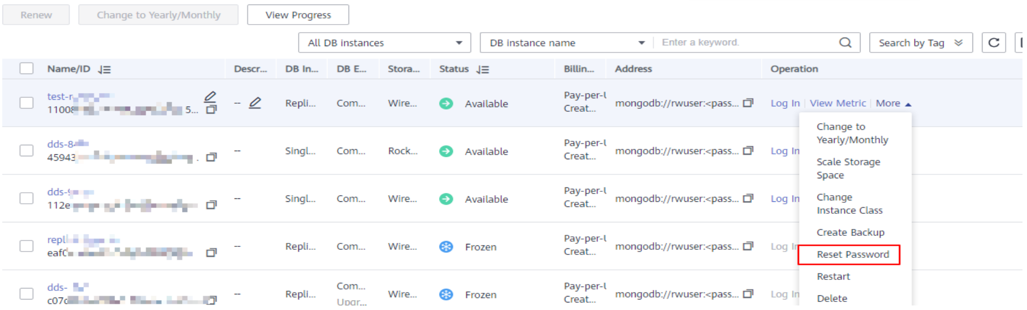

Name/ID	Descr...	DB In...	DB E...	Stora...	Status	Billin...	Address	Operation
test-r[redacted] 1100[redacted] 5...	--	Repli...	Com...	Wire...	Available	Pay-per-l Creat...	mongodb://rwuser:<pass...	Log In | View Metric | More
dds-8[redacted] 45943[redacted]	--	Singl...	Com...	Rock...	Available	Pay-per-l Creat...	mongodb://rwuser:<pass...	Log In
dds-[redacted] 112e[redacted]	--	Singl...	Com...	Wire...	Available	Pay-per-l Creat...	mongodb://rwuser:<pass...	Log In
repl[redacted] eafC[redacted]	--	Repli...	Com...	Wire...	Frozen	Pay-per-l Creat...	mongodb://rwuser:<pass...	Log In
dds-[redacted] c07[redacted]	--	Repli...	Com... Upgr...	Wire...	Frozen	Pay-per-l Creat...	mongodb://rwuser:<pass...	Log In

Alternatively, on the **Instances** page, click the instance. In the **DB Information** area on the **Basic Information** page, click **Reset Password** in the **Administrator** field.

Figure 13-8 Resetting a password

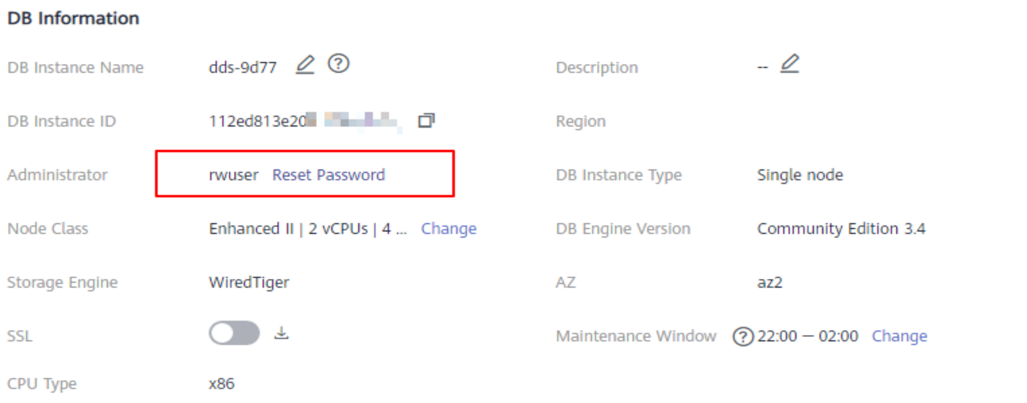

DB Information

DB Instance Name	dds-9d77	Description	--
DB Instance ID	112ed813e20[redacted]	Region	
Administrator	rwuser Reset Password	DB Instance Type	Single node
Node Class	Enhanced II | 2 vCPUs | 4 ... Change	DB Engine Version	Community Edition 3.4
Storage Engine	WiredTiger	AZ	az2
SSL		Maintenance Window	22:00 — 02:00 Change
CPU Type	x86		

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

- Resetting the password does not disconnect the authenticated connection. However, you will need to enter the new password when logging in to the database.
- The password must be 8 to 32 characters in length and contain uppercase letters, lowercase letters, digits, and any of the following special characters ~!@#%^*-_=+?()$

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

----End

13.3 Changing a Security Group

This section describes how to change a security group for cluster and replica set instances

Precautions

If any of the following operations is in progress, do not change the security group:

- Adding nodes
- Migrating data

Changing a Security Group

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the target DB instance.

Step 5 In the navigation pane on the left, choose **Connections**.

Step 6 In the **Security Group** area, click ✎ to select the security group to which the DB instance belongs.

Figure 13-9 Changing a security group

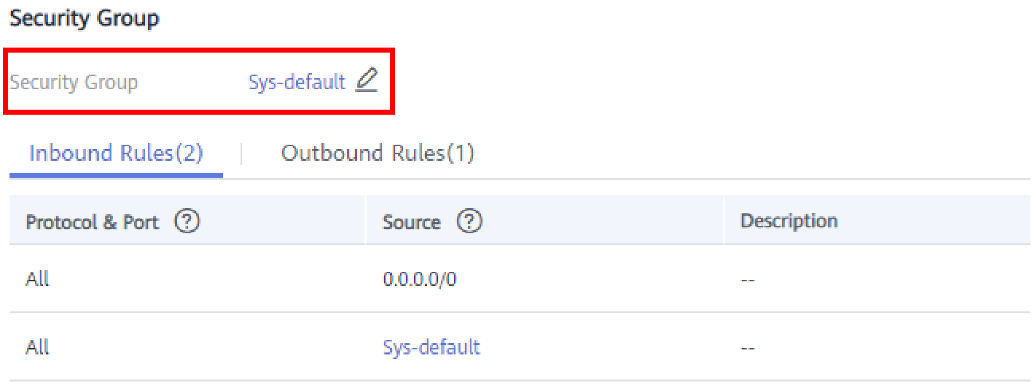

Security Group

Security Group Sys-default ✎

Inbound Rules(2) | Outbound Rules(1)

Protocol & Port ?	Source ?	Description
All	0.0.0.0/0	--
All	Sys-default	--

- To submit the change, click ✓. This process takes about 1 to 3 minutes.
- To cancel the change, click ✗.

Step 7 View the modification result.

----End

Managing Security Groups

NOTE

To use multiple security groups for a DDS instance, choose **Service Tickets > Create Service Ticket** in the upper right corner of the management console to apply for the required permissions.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ≡ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the **Network Information** area on the **Basic Information** page, click **Manage** next to the **Security Group** field.

Alternatively, choose **Connections** in the navigation pane on the left. In the **Security Group** area, click **Manage**.

- You can select multiple security groups at a time. The security group rules will be applied based on the following sequence: the first security group associated will take precedence over those associated later, then the rule with the highest priority in that security group will be applied first.
- To create a new security group, click **Create Security Group**.

NOTE

Using multiple security groups may impact the network performance. Selecting more than five security groups is not recommended.

Figure 13-10 Managing security groups

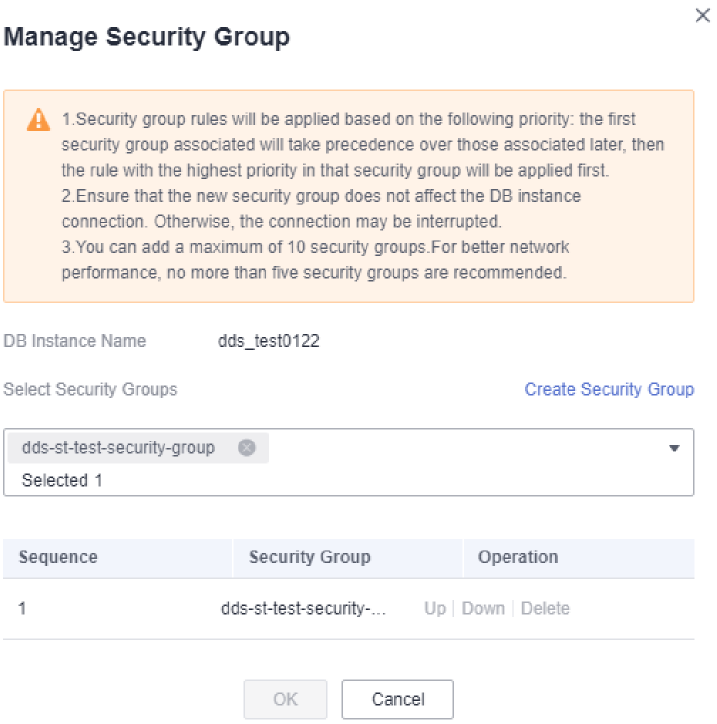

Manage Security Group

1.Security group rules will be applied based on the following priority: the first security group associated will take precedence over those associated later, then the rule with the highest priority in that security group will be applied first.
2.Ensure that the new security group does not affect the DB instance connection. Otherwise, the connection may be interrupted.
3.You can add a maximum of 10 security groups.For better network performance, no more than five security groups are recommended.

DB Instance Name dds_test0122

Select Security Groups

dds-st-test-security-group
Selected 1

Sequence	Security Group	Operation
1	dds-st-test-security-...	Up | Down | Delete

Step 6 Click **OK**.

----End

14 Monitoring and Alarm Reporting

14.1 DDS Metrics

This section describes metrics reported by Document Database Service (DDS) to Cloud Eye as well as their namespaces and dimensions. You can use APIs provided by Cloud Eye to query the metrics of the monitored object and alarms generated for DDS.

Namespace

SYS.DDS

Monitoring Metrics

Table 14-1 Recommended DDS metrics

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo007_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	0~100%	• dds mongos node • Primary node • Secondary node	1 minute 5 seconds

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo032_mem_usage	Memory Usage	Memory usage of the monitored object	0~100%	● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo031_cpu_usage	CPU Usage	CPU usage of the monitored object	0~100%	● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo035_disk_usage	Storage Space Usage	Storage space usage of the monitored object	0~100%	● Primary node ● Secondary node	1 minute

Table 14-2 DDS metrics

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo001_command_ps	COMMAND Statements per Second	Number of COMMAND statements executed per second	≥ 0 Count/s	● DDS DB instance ● dds mongos node ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute 5 seconds

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo002_delete_ps	DELETE Statements per Second	Number of DELETE statements executed per second	≥ 0 Count/s	● DDS DB instance ● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo003_insert_ps	INSERT Statements per Second	Number of INSERT statements executed per second	≥ 0 Count/s	● DDS DB instance ● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo004_query_ps	QUERY Statements per Second	Number of QUERY statements executed per second	≥ 0 Count/s	● DDS DB instance ● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo005_update_ps	UPDATE Statements per Second	Number of UPDATE statements executed per second	≥ 0 Count/s	● DDS DB instance ● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo006_getmore_ps	GETMORE Statements per Second	Number of GETMORE statements executed per second	≥ 0 Count/s	● DDS DB instance ● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo007_chunk_num1	Chunks of Shard 1	Number of chunks in shard 1	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num2	Chunks of Shard 2	Number of chunks in shard 2	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num3	Chunks of Shard 3	Number of chunks in shard 3	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num4	Chunks of Shard 4	Number of chunks in shard 4	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num5	Chunks of Shard 5	Number of chunks in shard 5	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num6	Chunks of Shard 6	Number of chunks in shard 6	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num7	Chunks of Shard 7	Number of chunks in shard 7	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num8	Chunks of Shard 8	Number of chunks in shard 8	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num9	Chunks of Shard 9	Number of chunks in shard 9	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num10	Chunks of Shard 10	Number of chunks in shard 10	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num11	Chunks of Shard 11	Number of chunks in shard 11	0–64 Counts	DDS cluster instance	1 minute
mongo007_chunk_num12	Chunks of Shard 12	Number of chunks in shard 12	0–64 Counts	DDS cluster instance	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo008_connections	Active Instance Connections	Total number of connections attempting to connect to a DDS DB instance	0–200 Counts	DDS DB instance	1 minute
mongo009_migFail_num	Chunk Migration Failures in Last 24 hrs	Number of chunk migration failures in the last 24 hours	≥ 0 Counts	DDS cluster instance	1 minute
mongo007_connections	Active Node Connections	Total number of connections attempting to connect to a DDS DB instance node	0–200 Counts	● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo007_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	0~100%	● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo008_mem_resident	Resident Memory	Size of resident memory in MB	≥ 0 MB	● dds mongos node ● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo009_mem_virtual	Virtual Memory	Size of virtual memory in MB	≥ 0 MB	● dds mongos node ● Primary node ● Secondary node	1 minute
mongo010_regular_asserts_ps	Regular Asserts per Second	Number of regular asserts per second	≥ 0 Count/s	● dds mongos node ● Primary node ● Secondary node	1 minute
mongo011_warning_asserts_ps	Warning Asserts per Second	Number of warning asserts per second	≥ 0 Count/s	● dds mongos node ● Primary node ● Secondary node	1 minute
mongo012_msg_asserts_ps	Message Asserts per Second	Number of message asserts per second	≥ 0 Count/s	● dds mongos node ● Primary node ● Secondary node	1 minute
mongo013_user_asserts_ps	User Asserts per Second	Number of user asserts per second	≥ 0 Count/s	● dds mongos node ● Primary node ● Secondary node	1 minute
mongo014_queues_total	Operations Queued Waiting for a Lock	Number of operations queued waiting for a lock	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo015_queues_readers	Operations Queued Waiting for a Read Lock	Number of operations queued waiting for a read lock	≥ 0 Counts	● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo016_queues_writers	Operations Queued Waiting for a Write Lock	Number of operations queued waiting for a write lock	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo017_page_faults	Page Faults	Number of page faults on the monitored nodes	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo018_porfling_num	Slow Queries	Total number of slow queries from the last 5 minutes to the current time point on the monitored node.	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo019_cursors_open	Maintained Cursors	Number of maintained cursors on the monitored nodes	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo020_cursors_timeOut	Timeout Cursors	Number of timed out cursors on the monitored nodes	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo021_wt_cahe_usage	Bytes in WiredTiger Cache	Size of data in the WiredTiger cache in MB	≥ 0 MB	● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo022_wt_cahe_dirty	Tracked Dirty Bytes in WiredTiger Cache	Size of tracked dirty data in the WiredTiger cache in MB	≥ 0 MB	● Primary node ● Secondary node	1 minute
mongo023_wInto_wtCache	Bytes Written Into Cache per Second	Bytes written into WiredTiger cache per second	≥ 0 bytes/s	● Primary node ● Secondary node	1 minute
mongo024_wFrom_wtCache	Bytes Written From Cache per Second	Bytes written from the WiredTiger cache to the disk per second	≥ 0 bytes/s	● Primary node ● Secondary node	1 minute
mongo025_repl_oplog_win	Oplog Window	Available time in hour in the monitored primary node's oplog	≥ 0 Hours	Primary node	1 minute
mongo025_repl_headroom	Replication Headroom	Time difference in seconds between the primary's oplog window and the replication lag of the secondary	≥ 0 Seconds	Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo026_repl_lag	Replication Lag	A delay in seconds between an operation on the primary and the application of that operation from the oplog to the secondary	≥ 0 Seconds	Secondary node	1 minute
mongo027_repl_command_ps	Replicated COMMAND Statements per Second	Number of replicated COMMAND statements executed on the secondary node per second	≥ 0 Count/s	Secondary node	1 minute
mongo028_repl_update_ps	Replicated UPDATE Statements per Second	Number of replicated UPDATE statements executed on the secondary node per second	≥ 0 Count/s	Secondary node	1 minute
mongo029_repl_delete_ps	Replicated DELETE Statements per Second	Number of replicated DELETE statements executed on the secondary node per second	≥ 0 Count/s	Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo030_repl_insert_ps	Replicated INSERT Statements per Second	Number of replicated INSERT statements executed on the secondary node per second	≥ 0 Count/s	Secondary node	1 minute
mongo031_cpu_usage	CPU Usage	CPU usage of the monitored object	0~100%	● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo032_mem_usage	Memory Usage	Memory usage of the monitored object	0~100%	● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo033_bytes_out	Network Output Throughput	Outgoing traffic in bytes per second	≥ 0 bytes/s	● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo034_bytes_in	Network Input Throughput	Incoming traffic in bytes per second	≥ 0 bytes/s	● dds mongos node ● Primary node ● Secondary node	1 minute 5 seconds
mongo035_disk_usage	Storage Space Usage	Storage space usage of the monitored object	0~100%	● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo036_iops	IOPS	Average number of I/O requests processed by the system in a specified period	≥ 0 Count/s	● Primary node ● Secondary node	1 minute
mongo037_read_throughput	Disk Read Throughput	Number of bytes read from the disk per second	≥ 0 bytes/s	● Primary node ● Secondary node	1 minute
mongo038_write_throughput	Disk Write Throughput	Number of bytes written into the disk per second	≥ 0 bytes/s	● Primary node ● Secondary node	1 minute
mongo039_avg_disk_sec_per_read	Average Time per Disk Read	Average time required for each disk read in a specified period	≥ 0 Seconds	● Primary node ● Secondary node	1 minute
mongo040_avg_disk_sec_per_write	Average Time per Disk Write	Average time required for each disk write in a specified period	≥ 0 Seconds	● Primary node ● Secondary node	1 minute
mongo042_disk_total_size	Total Storage Space	Total storage space of the monitored object	0–1000 GB	● Primary node ● Secondary node	1 minute
mongo043_disk_used_size	Used Storage Space	Used storage space of the monitored object	0–1000 GB	● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo044_swap_usage	SWAP Usage	Swap usage, in percentage.	0~100%	● dds mongos node ● Secondary node	1 minute
mongo050_top_total_time	Total Time Spent on Collections	Mongotop-total time: total time spent on collection operations, in milliseconds.	≥ 0 Milliseconds	● Primary node ● Secondary node	1 minute
mongo051_top_read_time	Total Time Spent on Collections	Mongotop-read time: total time spent reading collections, in milliseconds.	≥ 0 Milliseconds	● Primary node ● Secondary node	1 minute
mongo052_top_write_time	Total Time Spent on Collections	Mongotop-write time: total time spent writing collections, in milliseconds.	≥ 0 Milliseconds	● Primary node ● Secondary node	1 minute
mongo053_wt_flushes_status	Number of Times that Checkpoints Are Triggered	Number of times that the checkpoint is triggered during a polling interval of WiredTiger	≥ 0 Counts	● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo054_wt_cache_used_percent	Percentage of the Cache Used by WiredTiger	Cache size used by WiredTiger, in percentage	0~100%	● Primary node ● Secondary node	1 minute
mongo055_wt_cache_dirty_percent	Percentage of Dirty Data in the WiredTiger Cache	Dirty size in the WiredTiger cache, in percentage	0~100%	● Primary node ● Secondary node	1 minute
mongo070_rocks_active_memtable	Memtable Data Size	Size of data in the active memtable	0~100 GB	● Primary node ● Secondary node	1 minute
mongo071_rocks_oplogcf_active_memtable	Memtable Data Size on Oplogcf	Size of data in the active memtable on oplogcf	0~100 GB	● Primary node ● Secondary node	1 minute
mongo072_rocks_all_memtable	Total Data Size of Memtable and Immutable-memtable	Total data size of memtable and immutable-memtable	0~100 GB	● Primary node ● Secondary node	1 minute
mongo073_rocks_oplogcf_all_memtable	Total Data Size of Memtable and Immutable-memtable on Oplogcf	Total data size of memtable and immutable-memtable on oplogcf	0~100 GB	● Primary node ● Secondary node	1 minute
mongo074_rocks_snapshots	Unreleased Snapshots	Number of unreleased snapshots	≥ 0 Counts	● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo075_rocks_oplogcf_snapshots	Unreleased Snapshots on Oplogcf	Number of unreleased snapshots on oplogcf	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo076_rocks_live_versions	Active Versions	Number of active versions	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo077_rocks_oplogcf_live_versions	Active Versions on Oplogcf	Number of active versions on oplogcf	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo078_rocks_block_cache	Data Size in Blockcache	Size of data in blockcache	0~100 GB	● Primary node ● Secondary node	1 minute
mongo079_rocks_background_errors	Accumulated Background Errors	Accumulated number of background errors	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo080_rocks_oplogcf_background_errors	Accumulated Background Errors on Oplogcf	Number of accumulated background errors on oplogcf	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo081_rocks_conflict_bytes_usage	Buffer Usage for Processing Transaction Write Conflicts	Usage of the buffer for processing transaction write conflicts	0~100%	● Primary node ● Secondary node	1 minute
mongo082_rocks_uncommitted_keys	Uncommitted Keys	Number of uncommitted keys	≥ 0 Counts	● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo083_rocks_committed_keys	Committed Keys	Number of committed keys	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo084_rocks_alive_txn	Length of Active Transaction Linked Lists	Length of active transaction linked lists	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo085_rocks_read_queue	Length of Read Queues	Length of read queues	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo086_rocks_commit_queue	Length of Committed Queues	Length of committed queues	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo087_rocks_ct_write_out	Used Concurrent Write Transactions	Number of used concurrent write transactions	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo088_rocks_ct_write_available	Available Concurrent Write Transactions	Number of available concurrent write transactions	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo089_rocks_ct_read_out	Used Concurrent Read Transactions	Number of used concurrent read transactions	≥ 0 Counts	● Primary node ● Secondary node	1 minute
mongo090_rocks_ct_read_available	Available Concurrent Read Transactions	Number of available concurrent read transactions	≥ 0 Counts	● Primary node ● Secondary node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo091_active_session_count	Active Sessions	Number of active sessions cached in the memory of the Mongo instance since the last refresh	≥ 0 Counts	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden nodes of a DDS instance	1 minute
mongo092_rx_errors	Error Rate of Received Packets	Ratio of the number of error packets to the total number of received packets during the monitoring period	0~100%	DDS DB instance	1 minute 5 seconds
mongo093_rx_dropped	Loss Rate of Received Packets	Ratio of the number of lost packets to the total number of received packets during the monitoring period	0~100%	DDS DB instance	1 minute 5 seconds

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo094_tx_errors	Error Rate of Sent Packets	Ratio of the number of error packets to the total number of sent packets during the monitoring period	0~100%	DDS DB instance	1 minute 5 seconds
mongo095_tx_dropped	Loss Rate of Sent Packets	Ratio of the number of lost packets to the total number of sent packets during the monitoring period	0~100%	DDS DB instance	1 minute 5 seconds
mongo096_retrans_segs	Retransmitted Packets	The number of retransmitted packets during the monitoring period	≥ 0 Counts	DDS DB instance	1 minute 5 seconds
mongo097_retrans_rate	Retransmission Ratio	Ratio of retransmitted packets during the monitoring period	0~100%	DDS DB instance	1 minute 5 seconds
mongo098_out_rsts_nums	Sent RST Packets	The number of sent RST packets during the monitoring period	≥ 0 Counts	DDS DB instance	1 minute 5 seconds

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo099_read_time_average	Average Read Latency	Average read command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute
mongo100_read_time_p99	P99 Read Latency	P99 read command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute
mongo101_read_time_p999	P999 Read Latency	P999 read command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo102_write_time_average	Average Write Latency	Average write command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute
mongo103_write_time_p99	P99 Write Latency	P99 write command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute
mongo104_write_time_p999	P999 Write Latency	P999 write command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo105_command_time_average	Average Command Latency	Average command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute
mongo106_command_time_p99	P99 Command Latency	P99 command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute
mongo107_command_time_p999	P999 Command Latency	P999 command execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute

Metric ID	Metrics Name	Description	Value Range	Monitored Object	Monitoring Interval (Raw Data)
mongo108_txn_time_average	Average Transaction Latency	Average transaction execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute
mongo109_txn_time_p99	P99 Transaction Latency	P99 transaction execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute
mongo110_txn_time_p999	P999 Transaction Latency	P999 transaction execution latency of a single node	≥ 0 Milliseconds	● DDS DB instance ● Read replica of a DDS replica set instance ● Primary node ● Secondary node ● Hidden node	1 minute

NOTE

Metrics whose IDs contain rocks are used to monitor instances or instance nodes of version 4.2.

Dimensions

Key	Value
mongodb_instance_id	DDS DB instance ID Supports cluster instances of Community Edition, replica set instances, and single node instances.
mongodb_node_id	DDS node ID

NOTE

mongodb_instance_id is used to specify dimension fields when the Cloud Eye API is invoked. Replica sets and single node instance types do not have instance-level metrics.

14.2 Configuring Monitoring by Seconds

The default monitoring interval is 1 minute. To improve the instantaneous accuracy of monitoring metrics, you can set the monitoring interval to 5 seconds.

Precautions

- Only some monitoring metrics support monitoring by seconds. For details, see **Monitoring Metrics**.
- To apply for the advanced O&M permission, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console.
- To apply for the monitoring permission, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console.

Enabling Monitoring by Seconds

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the target instance name.

Step 5 In the navigation pane on the left, choose **Advanced O&M**.

Step 6 On the displayed page, click the **Real-Time Monitoring** tab and click ○ next to **Monitoring by Seconds**.

NOTICE

Instances with fewer than four vCPUs do not support monitoring by seconds.

Step 7 In the displayed dialog box, select a collection period and click **Yes**.

Monitoring by Seconds will be automatically disabled for instances with fewer than 4 vCPUs. After you enable this function, monitoring data will be reported again and will be displayed by seconds about five minutes later.

Figure 14-1 Enable Monitoring by Seconds

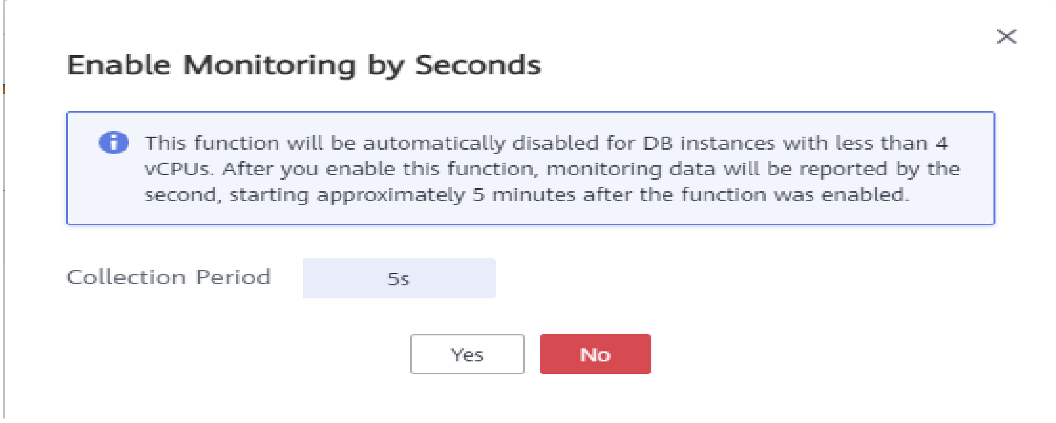

Enable Monitoring by Seconds

This function will be automatically disabled for DB instances with less than 4 vCPUs. After you enable this function, monitoring data will be reported by the second, starting approximately 5 minutes after the function was enabled.

Collection Period 5s

----End

Disabling Monitoring by Seconds

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the target instance name.

Step 5 In the navigation pane on the left, choose **Advanced O&M**.

Step 6 On the displayed page, click the **Real-Time Monitoring** tab and click next to **Monitoring by Seconds**.

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

After you disable this function, monitoring data will be reported again and will be displayed by the minute about five minutes later.

Figure 14-2 Disable Monitoring by Seconds

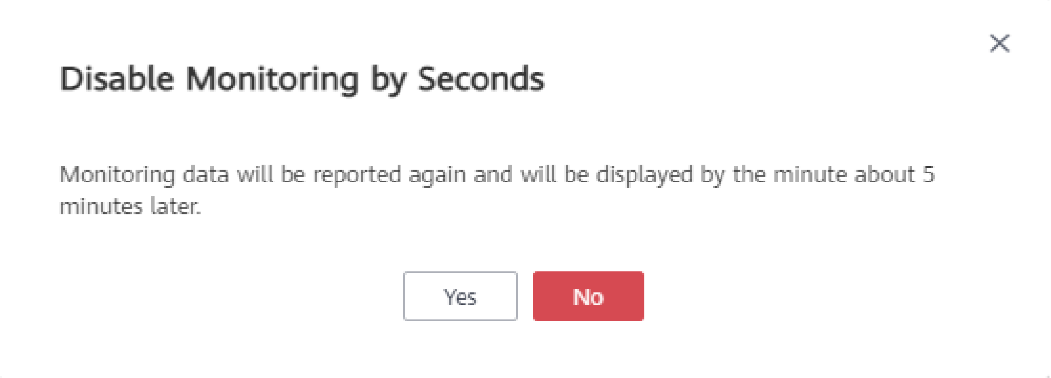

Disable Monitoring by Seconds

Monitoring data will be reported again and will be displayed by the minute about 5 minutes later.

----End

14.3 Viewing DDS Metrics

- Cloud Eye monitors DDS running statuses. You can obtain the monitoring metrics of DDS on the management console.
- Monitored data requires a period of time for transmission and display. The status of DDS displayed on the Cloud Eye page is from about 5 to 10 minutes ago, so the data for a newly created DB instance takes about 5 to 10 minutes to show up on Cloud Eye.
- The monitoring data is retained for 30 days.
- If you receive an alarm (for example, indicating that the data disk space is insufficient), you need to filter the instance nodes to check whether each node is normal when you view the instance monitoring data for problem location and analysis.

Prerequisites

- The DDS DB instance is running normally.

  Cloud Eye does not display the metrics of faulty or deleted DB instances or nodes. You can view the monitoring information only after the instance is restarted or recovered.
- The DB instance has been properly running for at least 10 minutes.

  For a newly created DB instance, you need to wait a bit before the monitoring metrics show up on Cloud Eye.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the target DB instance.

Step 5 In the navigation pane on the left, choose **Advanced O&M**.

Step 6 View metrics.

- For cluster instances, you can view metrics of instances, and dds mongos, shard, and config nodes.

Figure 14-3 Viewing metrics of a cluster instance

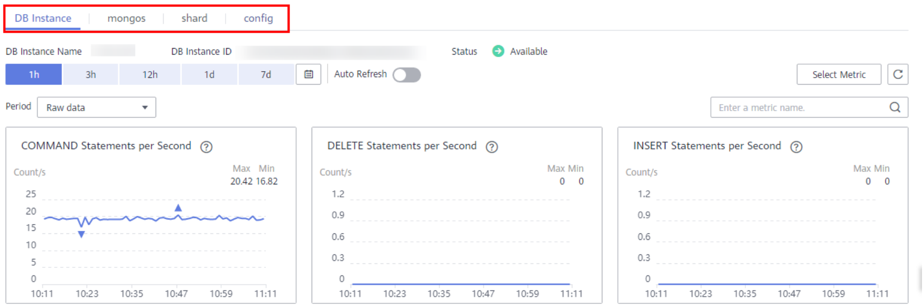

- For replica set instances, you can view metrics of primary, secondary, and hidden nodes.

Figure 14-4 Viewing metrics of a replica set instance

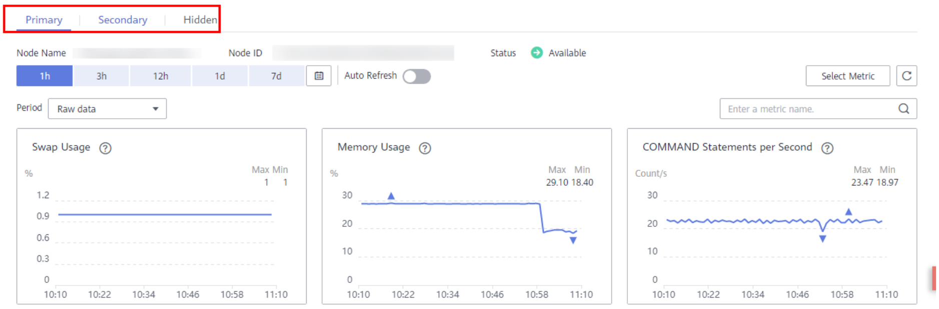

- For single node instances, you can view node metrics.

Figure 14-5 Viewing metrics of a single node instance

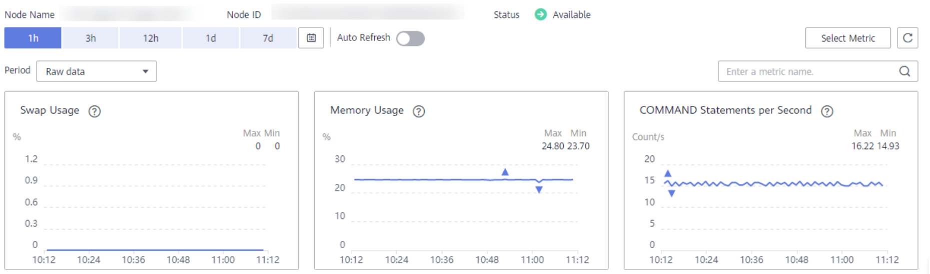

Step 7 View monitoring metrics of cluster instances, cluster instance nodes, and replica set instance nodes.

Step 8 In the DDS monitoring area, you can select a duration to view the monitoring data. You can view the monitoring data of the last 1 hour, 3 hours, and 12 hours.

Figure 14-6 Enabling Auto Refresh

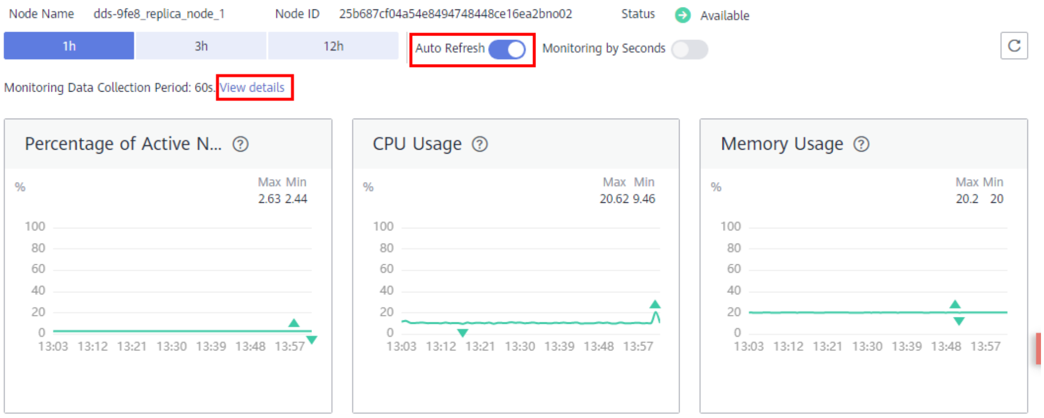

Node Name dds-9fe8_replica_node_1 Node ID 25b687cf04a54e8494748448ce16ea2bno02 Status Available

Auto Refresh: on. Monitoring by Seconds: off

Monitoring Data Collection Period: 60s. View details

- If automatic refresh is enabled, monitoring data is automatically refreshed every 60 seconds.
- For more metric information, click **View details** to switch to the Cloud Eye console.

----End

14.4 Configuring Alarm Rules

DDS allows you to set threshold rules for instance metrics. If the value of a metric exceeds the threshold, an alarm is triggered. The system automatically sends an alarm notification to the cloud account contact through SMN, helping you learn about the running status of the DDS instance in a timely manner.

You can configure alarm rules on the Cloud Eye console.

Precautions

The basic alarm function is free of charge. SMN sends you the alarm messages and charges you for that. For pricing details, see **Pricing Details**.

Customizing Alarm Rules

Step 1 Log in to the management console.

Step 2 Under **Management & Governance**, click **Cloud Eye**.

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

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

Step 5 On the **Create Alarm Rule** page, follow the prompts to set the parameters.

Pay attention to the following parameters:

- **Resource Type**: Select **Document Database Service**.
- **Dimension**: DDS supports instance-level and node-level monitoring dimensions. Different monitoring metrics support different monitoring dimensions. For details, see **DDS Metrics**.

Figure 14-7 Configuring monitoring dimensions

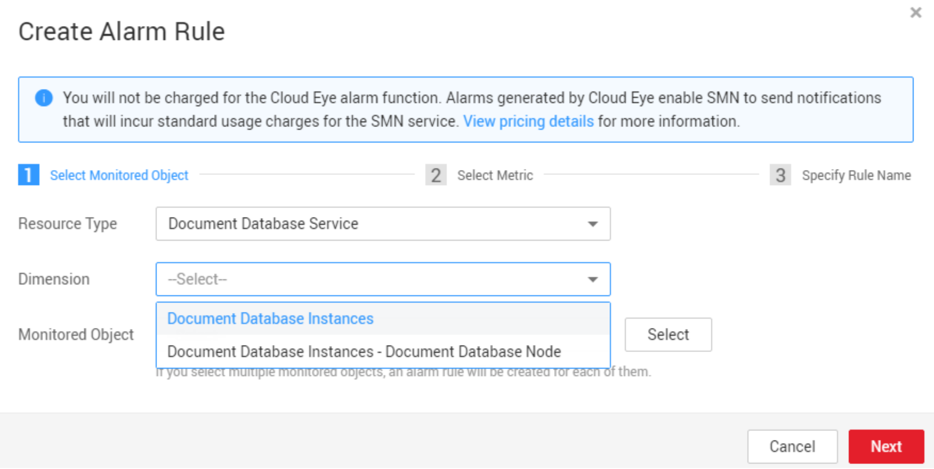

Create Alarm Rule

You will not be charged for the Cloud Eye alarm function. Alarms generated by Cloud Eye enable SMN to send notifications that will incur standard usage charges for the SMN service. View pricing details for more information.

1 Select Monitored Object — 2 Select Metric — 3 Specify Rule Name

Resource Type: Document Database Service

Dimension: --Select--

- Document Database Instances
- Document Database Instances - Document Database Node

Monitored Object: Select

If you select multiple monitored objects, an alarm rule will be created for each of them.

Step 6 After the alarm rule is set, the system automatically notifies you when an alarm is triggered.

----End

14.5 Managing Alarm Rules

This section describes how to enable and disable alarm reporting on the Cloud Eye console.

Disabling an Alarm Rule

Step 1 Log in to the management console.

Step 2 Under **Management & Governance**, click **Cloud Eye**.

Step 3 In the navigation pane on the left, choose **Alarm Management** > **Alarm Rules**, locate the alarm rule you want to disable and click **Disable** in the **Operation** column.

Figure 14-8 Disabling an Alarm Rule

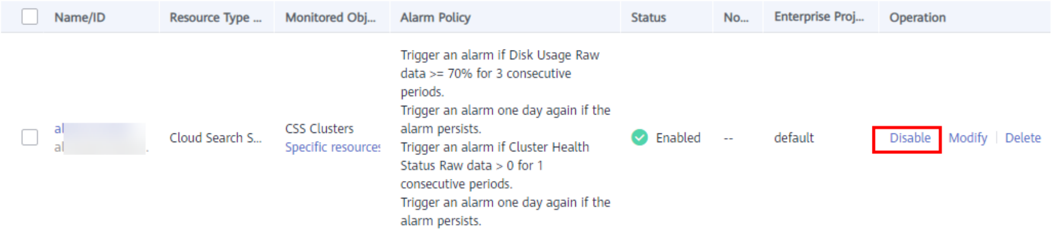

□ Name/ID	Resource Type ...	Monitored Obj...	Alarm Policy	Status	No...	Enterprise Proj...	Operation
□ al [redacted] al [redacted]	Cloud Search S...	CSS Clusters Specific resources	Trigger an alarm if Disk Usage Raw data >= 70% for 3 consecutive periods. Trigger an alarm one day again if the alarm persists. Trigger an alarm if Cluster Health Status Raw data > 0 for 1 consecutive periods. Trigger an alarm one day again if the alarm persists.	✓ Enabled	--	default	Disable | Modify | Delete

Step 4 In the displayed **Disable Alarm Rule** dialog box, click **Yes** to disable the alarm rule.

If you want to disable multiple alarm rules, on the **Alarm Rules** page, select multiple alarm rules, and click **Disable** in the upper left of the alarm rule list. In the displayed **Disable Alarm Rule** dialog box, click **Yes**.

----End

Enabling an Alarm Rule

Step 1 Log in to the management console.

Step 2 Under **Management & Governance**, click **Cloud Eye**.

Step 3 In the navigation pane on the left, choose **Alarm Management** > **Alarm Rules**, locate the alarm rule you want to enable and click **Enable** in the **Operation** column.

Figure 14-9 Enabling an Alarm Rule

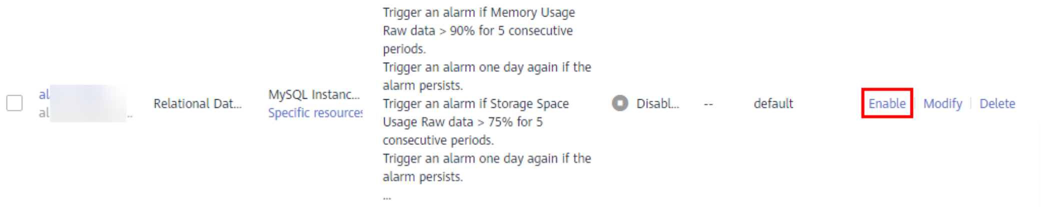

□	al al [redacted] ..	Relational Dat...	MySQL Instanc... Specific resources	Trigger an alarm if Memory Usage Raw data > 90% for 5 consecutive periods. Trigger an alarm one day again if the alarm persists. Trigger an alarm if Storage Space Usage Raw data > 75% for 5 consecutive periods. Trigger an alarm one day again if the alarm persists. ...	● Disabl...	--	default	Enable Modify | Delete

Step 4 In the displayed **Enable Alarm Rule** dialog box, click **Yes** to enable the alarm rule.

If you want to enable multiple alarm rules, on the **Alarm Rules** page, select multiple alarm rules, and click **Enable** in the upper left of the alarm rule list. In the displayed **Enable Alarm Rule** dialog box, click **Yes**.

----End

14.6 Event Monitoring

14.6.1 Introduction to Event Monitoring

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

Events are key operations on DDS that are stored and monitored by Cloud Eye. You can view events to see operations performed by specific users on specific resources, such as deleting a read replica or changing instance specifications.

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

Event monitoring is enabled by default. You can view monitoring details about system events and custom events. For details about system events, see **Events Supported by Event Monitoring**.

14.6.2 Viewing Event Monitoring Data

Scenarios

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

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

This topic describes how to view the event monitoring data.

Procedure

Step 1 Log in to the management console.

Step 2 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 3 On the **Instances** page, locate the instance and click **View Metric** in the **Operation** column to go to the Cloud Eye console.

Alternatively, go to the Cloud Eye console using either of the following methods:

- On the **Instances** page, locate the instance and click its name. On the displayed **Basic Information** page, click **View Metric** in the upper right corner to go to the Cloud Eye console.
- On the **Instances** page, locate the instance and click its name. On the displayed **Basic Information** page, in the **Node Information** area, click **View Metric** in the **Operation** column to go to the Cloud Eye console.
- On the **Instances** page, locate the instance and click its name. In the navigation pane on the left, click **Advanced O&M**, locate the target node, and click **View details** to go to the Cloud Eye console.

Step 4 Click < in the upper left corner to return to the main page of Cloud Eye.

Step 5 In the navigation pane on the left, choose **Event Monitoring**.

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

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

Step 6 Click ⌄ to expand an event, and click **View Event** in the **Operation** column to view details about a specific event.

----End

14.6.3 Creating an Alarm Rule to Monitor an Event

Scenarios

This topic describes how to create an alarm rule to monitor an event.

Procedure

Step 1 Log in to the management console.

Step 2 Click ☰ in the upper left corner of the page. Under **Management & Governance**, click **Cloud Eye**.

Step 3 In the navigation pane on the left, choose **Event Monitoring**.

Step 4 On the event list page, click **Create Alarm Rule** in the upper right corner.

Step 5 On the **Create Alarm Rule** page, configure the parameters.

Table 14-3 Parameter description

Parameter	Description
Name	Specifies the alarm rule name. The system generates a random name, which you can modify.
Description	(Optional) Provides supplementary information about the alarm rule.
Alarm Type	Specifies the alarm type corresponding to the alarm rule.
Event Type	Specifies the event type of the metric corresponding to the alarm rule.
Event Source	Specifies the service the event is generated for. Select **Document Database Service**.
Monitoring Scope	Specifies the monitoring scope for event monitoring.
Method	Select **Configure manually**.
Alarm Policy	**Event Name** indicates the instantaneous operations users performed on system resources, such as login and logout. For details about events supported by Event Monitoring, see **Events Supported by Event Monitoring**. Select **Trigger Mode** and **Alarm Severity** as required.

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

Table 14-4 Alarm notification parameters

Parameter	Description
Alarm Notification	Specifies whether to notify users when alarms are triggered. Notifications can be sent by email, text message, or HTTP/HTTPS message.
Notification Type	You can select a notification group or topic subscription as required.
Notification Group	Specifies the notification group that needs to send alarm notifications.

Parameter	Description
Notification Object	Specifies the object that receives alarm notifications. You can select the account contact or a topic. ● **Account contact** is the mobile phone number and email address of the registered account. ● **Topic** is used to publish messages and subscribe to notifications. If the required topic is unavailable, create one first and add subscriptions to it. For details, see **Creating a Topic** and **Adding Subscriptions**.
Validity Period	Cloud Eye sends notifications only within the validity period specified in the alarm rule. If **Validity Period** is set to **08:00-20:00**, Cloud Eye sends notifications only within 08:00-20:00.
Trigger Condition	Specifies the condition for triggering an alarm notification.

Configure the enterprise project as prompted.

Table 14-5 Parameter description

Parameter	Description
Enterprise Project	Specifies the enterprise project that the alarm rule belongs to. Only users with the enterprise project permissions can view and manage the alarm rule. For details about how to create an enterprise project, see **Creating an Enterprise Project**.

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

----End

14.6.4 Events Supported by Event Monitoring

Table 14-6 Document Database Service (DDS)

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
DDS	DB instance creation failure	DDSCreateInstanceFailed	Major	A DDS instance fails to be created due to insufficient disks, quotas, and underlying resources.	Check the number and quota of disks. Release resources and create DDS instances again.	DDS instances cannot be created.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Replication failed	DDSAbnormalReplicationStatus	Major	The possible causes are as follows: 1. The replication delay between the primary and secondary nodes is too long, which usually occurs when a large amount of data is written to databases or a large transaction is performed. During off-peak hours, the replication delay gradually decreases. 2. The network between the primary and secondary nodes is disconnected.	Submit a service ticket.	Your applications are not affected because this event does not interrupt data read and write.

Event Source	Event Name	Event ID	Event Severity	Description	Solution	Impact
	Replication recovered	DDSReplicationStatusRecovered	Major	The replication delay between the primary and standby instances is within the normal range, or the network connection between them has restored.	No further action is required.	None
	DB instance faulty	DDSFaultyDBInstance	Major	This event is a key alarm event and is reported when an instance is faulty due to a disaster or a server failure.	Submit a service ticket.	The database service may be unavailable.
	DB instance recovered	DDSDBInstanceRecovered	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
	Faulty node	DDSFaultyDBNode	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 available and submit a service ticket.	The database service may be unavailable.
	Node recovered	DDSDBNodeRecovered	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	DDSPrimaryStandbySwitched	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
	Insufficient storage space	DDSRiskyDataDiskUsage	Major	The storage space is insufficient.	Scale up storage space. For details, see section "Scaling Up Storage Space" in the corresponding user guide.	The instance is set to read-only and data cannot be written to the instance.
	Data disk expanded and being writable	DDSDataDiskUsageRecovered	Major	The capacity of a data disk has been expanded and the data disk becomes writable.	No action is required.	No adverse impact.
	Schedule for deleting a KMS key	DDSplanDeleteKmsKey	Major	A KMS key is scheduled to be deleted.	After a KMS key is scheduled to be deleted, either decrypt the data encrypted by KMS key in a timely manner or cancel the key deletion.	After a KMS key is deleted, disk encryption cannot be enabled.

15 Auditing

15.1 Key Operations Recorded by CTS

With Cloud Trace Service (CTS), you can record operations associated with DDS for later query, audit, and backtrack operations.

Table 15-1 Key operations on DDS

Operation	Resource	Trace Name
Restoring data to a new DB instance	instance	ddsRestoreToNewInstance
Restoring to an existing DB instance	instance	ddsRestoreToOldInstance
Creating a DB instance	instance	ddsCreateInstance
Deleting a DB instance	instance	ddsDeleteInstance
Restarting a DB instance	instance	ddsRestartInstance
Scaling up a DB instance	instance	ddsGrowInstance
Scaling up storage space	instance	ddsExtendInstanceVolume
Resetting the database password	instance	ddsResetPassword
Renaming a DB instance	instance	ddsRenameInstance
Switching SSL	instance	ddsSwitchSsl
Modifying a DB instance port	instance	ddsModifyInstancePort
Creating a backup	backup	ddsCreateBackup
Deleting a backup	backup	ddsDeleteBackup

Operation	Resource	Trace Name
Setting a backup policy	backup	ddsSetBackupPolicy
Applying a parameter template	parameterGroup	ddsApplyConfigurations
Replicating a parameter template	parameterGroup	ddsCopyConfigurations
Resetting a parameter template	parameterGroup	ddsResetConfigurations
Creating a parameter template	parameterGroup	ddsCreateConfigurations
Deleting a parameter template	parameterGroup	ddsDeleteConfigurations
Updating a parameter template	parameterGroup	ddsUpdateConfigurations
Binding an EIP	instance	ddsBindEIP
Unbinding an EIP	instance	ddsUnBindEIP
Editing a tag	tag	ddsModifyTag
Deleting an instance tag	tag	ddsDeleteInstanceTag
Adding an instance tag	tag	ddsAddInstanceTag
Rolling back upon scaling-up failure	instance	ddsDeleteExtendedDdsNode
Changing DB instance classes	instance	ddsResizeInstance
Unfreezing a DB instance	instance	ddsUnfreezeInstance
Freezing a DB instance	instance	ddsFreezeInstance
Changing a private IP address	instance	ddsModifyIP
Modifying a private domain name	instance	ddsModifyDNSName
Enabling or disabling cluster balancing	instance	ddsSetBalancer
Switching the internal communication mode	instance	ddsSwitchInnerSsl
Adding read replicas	instance	AddReadonlyNode
Enabling shard/config IP address for a cluster instance	instance	ddsCreateIp

Operation	Resource	Trace Name
Changing a security group	instance	ddsModifySecurityGroup
Changing an AZ	instance	ddsMigrateAvailabilityZone
Modifying instance remarks	instance	ddsModifyInstanceRemark
Configuring a maintenance window	instance	ddsModifyInstanceMaintenanceWindow
Upgrading patches	instance	ddsUpgradeDatastorePatch
Performing a primary/ standby switchover	instance	ddsReplicaSetSwitchover
Configuring cross-CIDR access	instance	ddsModifyInstanceSourceSubnet
Modifying instance parameters	parameterGroup	ddsUpdateInstanceConfigurations
Exporting a parameter template for a DB instance	parameterGroup	ddsSaveConfigurations
Setting a cross-region backup policy	backup	ddsModifyOffsiteBackupPolicy
Enabling plaintext display of slow query logs	instance	ddsOpenSlowLogPlaintextSwitch
Disabling plaintext display of slow query logs	instance	ddsCloseSlowLogPlaintextSwitch
Downloading error or slow query logs	instance	ddsDownloadLog
Enabling the audit policy for a DB instance	instance	ddsOpenAuditLog
Disabling the audit policy for a DB instance	instance	ddsCloseAuditLog
Downloading audit logs for a DB instance	instance	ddsDownloadAuditLog
Deleting audit logs for a DB instance	instance	ddsDeleteAuditLogFile
Modifying recycling policy	instance	ddsModifyRecyclePolicy

15.2 Viewing CTS Traces

For details about how to view audit logs, see **Querying Real-Time Traces**.

16 Logs

16.1 Log Reporting

Prerequisites

You have created a log group and a log stream on the Log Tank Service (LTS) console.

Scenarios

If you enable log reporting to LTS, new audit logs, error logs, and slow query logs generated for DDS DB instances will be uploaded to LTS for management. You can view details about audit logs, error logs, and slow query logs of DDS DB instances, including searching for logs, visualizing logs, downloading logs, and viewing real-time logs.

The following operations use audit logs as an example:

- Enable log reporting to LTS for a single DB instance by referring to **Enabling Log Reporting to LTS for a Single DB Instance**.
- Edit log reporting to LTS for a single DB instance by referring to **Editing Log Reporting to LTS for a Single DB Instance**.
- Disable log reporting to LTS for a single DB instance by referring to **Disabling Log Reporting to LTS for a Single DB Instance**.
- Enable log reporting to LTS in batches by referring to **Enabling Log Reporting to LTS in Batches**.
- Disable log reporting to LTS in batches by referring to **Disabling Log Reporting to LTS in Batches**.

Precautions

- To apply for the log reporting permission, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console.
- Logs record all requests sent to your DB instance and are stored in LTS.

- This request does not take effect immediately. There is a delay of about 10 minutes.
- You will be billed for log reporting. For details, see **LTS Pricing Details**.
- After this function is enabled, all audit policies are reported by default.
- Audit logs are generated every hour. If the size of an audit log exceeds 10 MB, a new audit log is generated.
- If **Audit Policy** is enabled, LTS reuses the audit policy set for your DB instance and you will also be billed for reporting audit logs to LTS. (Only after you disable **Audit Policy**, the fee will be terminated.)
- If you enable audit log reporting to LTS for an instance with the **Audit Policy** toggle switch turned on, you can turn off this switch only when the instance status becomes available.

Enabling Log Reporting to LTS for a Single DB Instance

Step 1 Log in to the console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate a target DB instance and click its name.

Step 5 In the navigation pane on the left, choose **Audit Logs**.

Step 6 On the **Audit Logs** page, click ○ next to **Report Logs to LTS**.

Figure 16-1 Enabling Report Logs to LTS

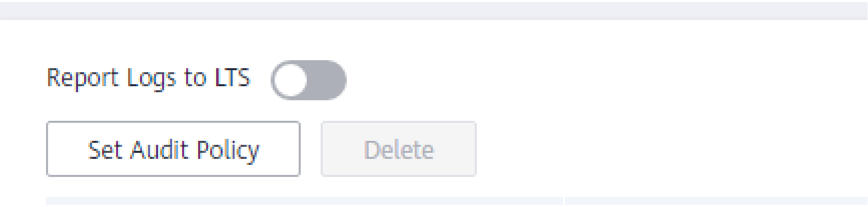

Step 7 In the displayed dialog box, specify **Log Group** and **Log Stream**.

Figure 16-2 Enabling audit log reporting to LTS

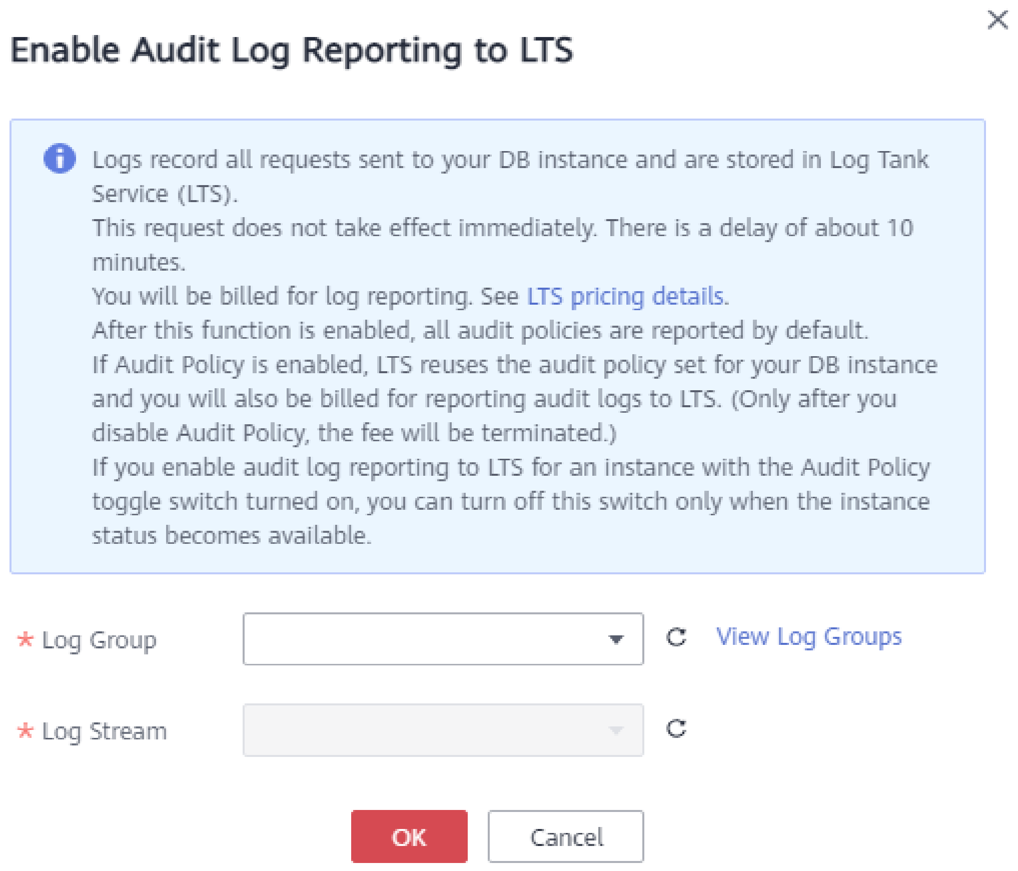

Enable Audit Log Reporting to LTS

Logs record all requests sent to your DB instance and are stored in Log Tank Service (LTS).
This request does not take effect immediately. There is a delay of about 10 minutes.
You will be billed for log reporting. See LTS pricing details.
After this function is enabled, all audit policies are reported by default.
If Audit Policy is enabled, LTS reuses the audit policy set for your DB instance and you will also be billed for reporting audit logs to LTS. (Only after you disable Audit Policy, the fee will be terminated.)
If you enable audit log reporting to LTS for an instance with the Audit Policy toggle switch turned on, you can turn off this switch only when the instance status becomes available.

NOTE

If you enable log reporting to LTS for the first time, click **View Log Groups** to log in to the LTS console and configure log groups and log streams. For details, see **Managing Log Groups** and **Managing Log Streams**.

Step 8 Click **OK**.

----End

Editing Log Reporting to LTS for a Single DB Instance

Step 1 Log in to the console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate a target DB instance and click its name.

Step 5 In the navigation pane on the left, choose **Audit Logs**.

Step 6 On the **Audit Logs** page, click **Edit** next to the **Report Logs to LTS** toggle switch.

NOTE

The editing function is available only when the **Report Logs to LTS** toggle switch is turned on.

Step 7 In the displayed dialog box, specify **Log Group** and **Log Stream**.

NOTE

Select the target log group and log stream.

Figure 16-3 Editing audit log reporting to LTS

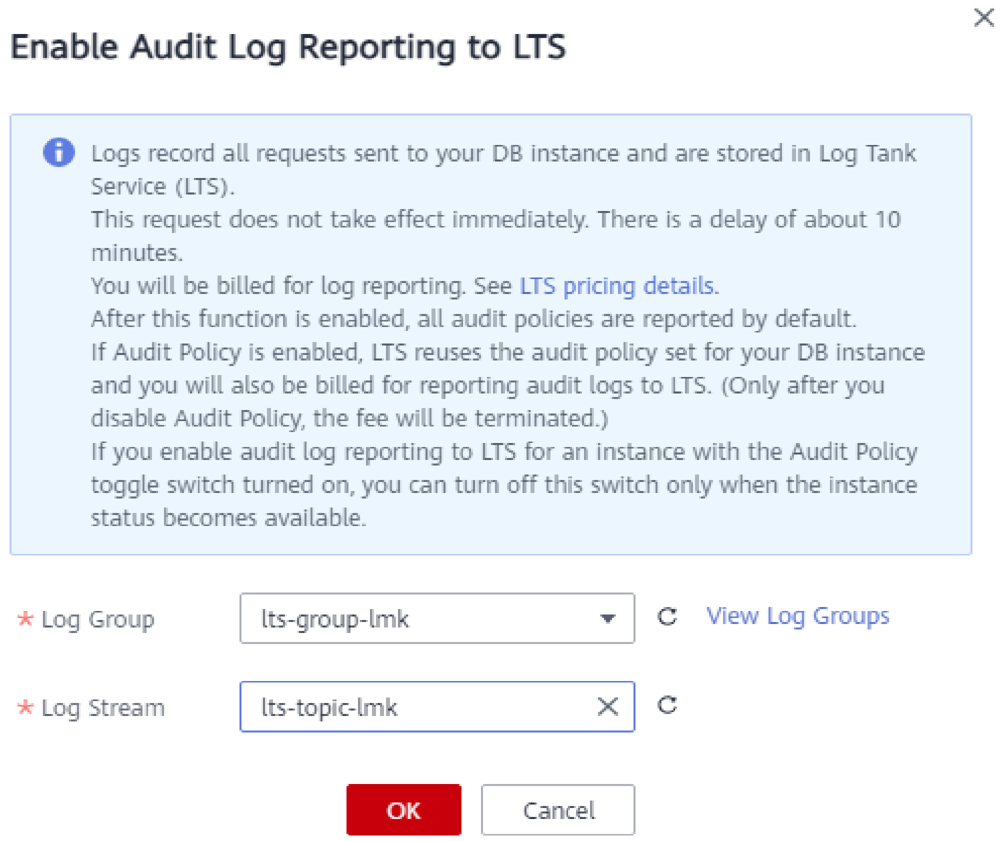

Enable Audit Log Reporting to LTS

Logs record all requests sent to your DB instance and are stored in Log Tank Service (LTS).
This request does not take effect immediately. There is a delay of about 10 minutes.
You will be billed for log reporting. See LTS pricing details.
After this function is enabled, all audit policies are reported by default.
If Audit Policy is enabled, LTS reuses the audit policy set for your DB instance and you will also be billed for reporting audit logs to LTS. (Only after you disable Audit Policy, the fee will be terminated.)
If you enable audit log reporting to LTS for an instance with the Audit Policy toggle switch turned on, you can turn off this switch only when the instance status becomes available.

* Log Group: lts-group-lmk

* Log Stream: lts-topic-lmk

Step 8 Click **OK**.

----End

Disabling Log Reporting to LTS for a Single DB Instance

Step 1 Log in to the console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate a target DB instance and click its name.

Step 5 In the navigation pane on the left, choose **Audit Logs**.

Step 6 On the **Audit Logs** page, click next to **Report Logs to LTS**.

Figure 16-4 Disabling Report Logs to LTS

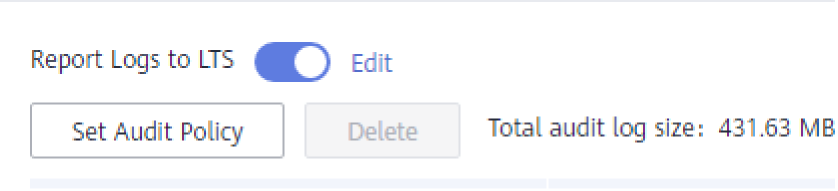

Report Logs to LTS Edit

Set Audit Policy | Delete | Total audit log size: 431.63 MB

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

Figure 16-5 Disabling audit log reporting to LTS

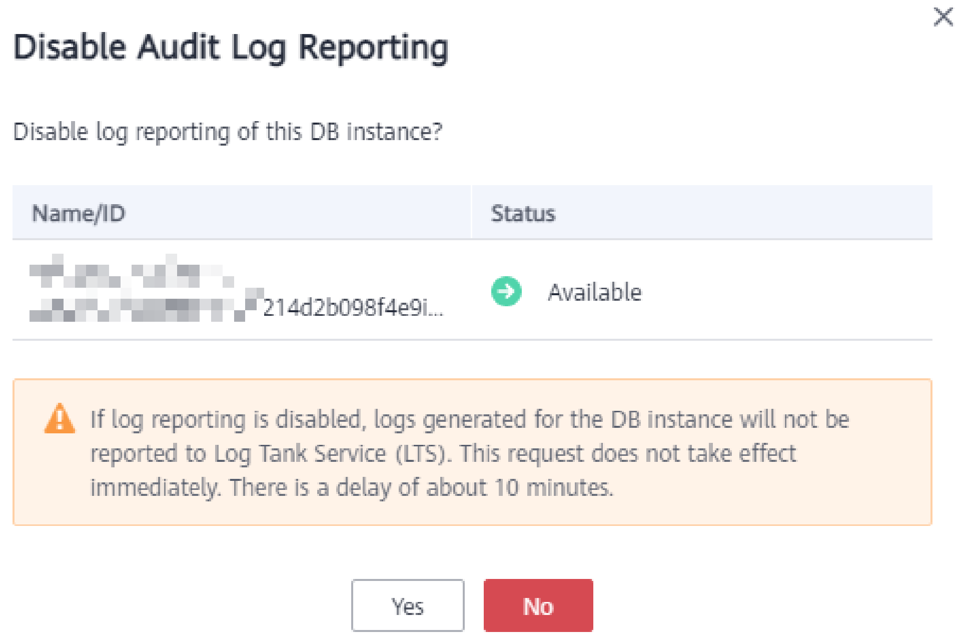

Disable Audit Log Reporting

Disable log reporting of this DB instance?

Name/ID	Status
[redacted] 214d2b098f4e9i...	Available

If log reporting is disabled, logs generated for the DB instance will not be reported to Log Tank Service (LTS). This request does not take effect immediately. There is a delay of about 10 minutes.

----End

Enabling Log Reporting to LTS in Batches

Step 1 Log in to the console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Log Reporting**.

Step 5 Select target DB instances and click **Enable Log Reporting**.

Step 6 In the displayed dialog box, specify **Log Group** and **Log Stream**.

Figure 16-6 Enabling log reporting to LTS in batches

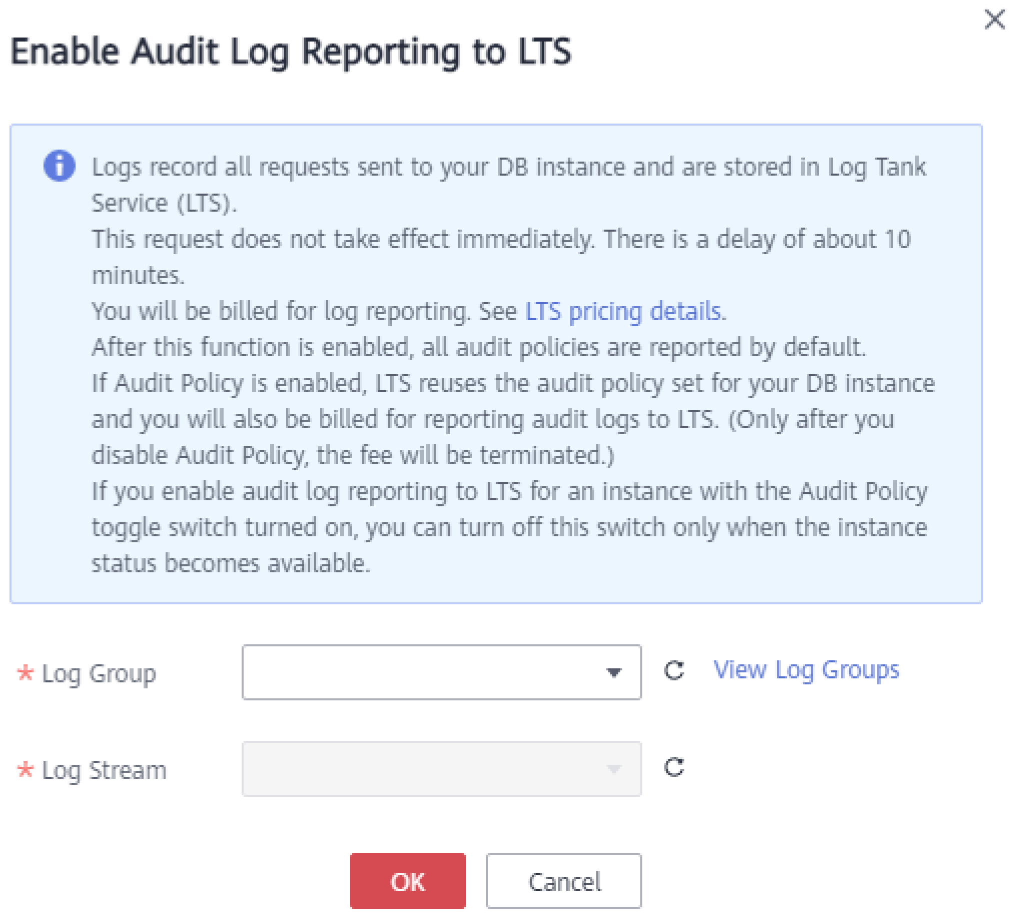

Enable Audit Log Reporting to LTS

Logs record all requests sent to your DB instance and are stored in Log Tank Service (LTS).
This request does not take effect immediately. There is a delay of about 10 minutes.
You will be billed for log reporting. See LTS pricing details.
After this function is enabled, all audit policies are reported by default.
If Audit Policy is enabled, LTS reuses the audit policy set for your DB instance and you will also be billed for reporting audit logs to LTS. (Only after you disable Audit Policy, the fee will be terminated.)
If you enable audit log reporting to LTS for an instance with the Audit Policy toggle switch turned on, you can turn off this switch only when the instance status becomes available.

NOTE

- Select the target log group and log stream.
- If you enable log reporting to LTS for the first time, click **View Log Groups** to log in to the LTS console and configure log groups and log streams. For details, see **Managing Log Groups** and **Managing Log Streams**.

Step 7 Click **OK**.

----End

Disabling Log Reporting to LTS in Batches

Step 1 Log in to the console.

Step 2 Click ⊙ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, choose **Log Reporting**.

Step 5 Select target DB instances and click **Disable Log Reporting**.

Figure 16-7 Disabling log reporting to LTS in batches

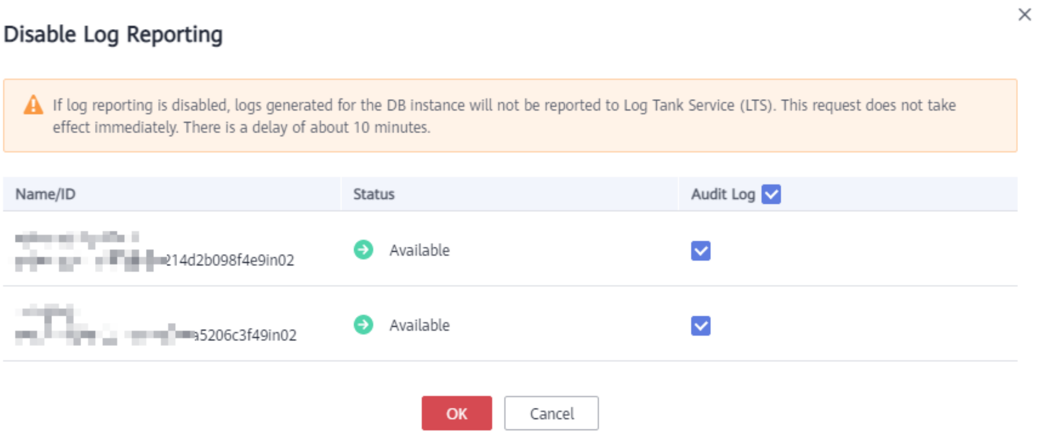

Disable Log Reporting

If log reporting is disabled, logs generated for the DB instance will not be reported to Log Tank Service (LTS). This request does not take effect immediately. There is a delay of about 10 minutes.

Name/ID	Status	Audit Log ☑
[redacted]214d2b098f4e9in02	Available	☑
[redacted]a5206c3f49in02	Available	☑

Step 6 In the displayed dialog box, click **OK**.

----End

16.2 Error Logs

16.2.1 Viewing Error Logs on the LTS Console

You can analyze, search for, monitor, download, and view real-time logs on the LTS console.

Querying Error Logs Reported to LTS

NOTE

You have enabled log reporting to LTS. For details, see **Log Reporting**.

Step 1 Click ☰ in the upper left corner of the page and choose **Management & Governance** > **Log Tank Service**.

Step 2 In the **Log Groups** area, locate a target log group and click its name. For details about logs, see **Log Management**.

Figure 16-8 Viewing log details

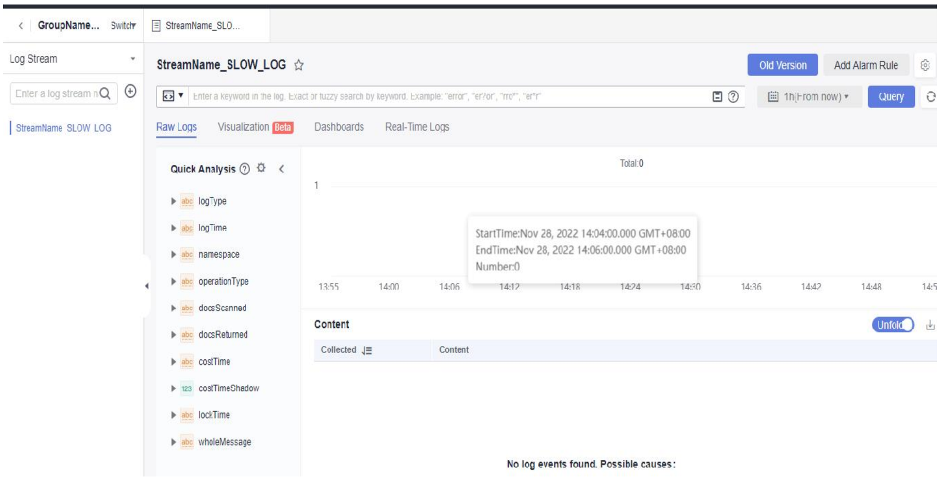

StartTime:Nov 28, 2022 14:04:00.000 GMT+08:00
EndTime:Nov 28, 2022 14:06:00.000 GMT+08:00
Number:0

No log events found. Possible causes:

----End

Downloading Error Logs Reported to LTS

NOTE

If you have enabled log reporting to LTS for your DB instance in **Log Reporting**, you can download logs on the LTS console.

Step 1 Click ☰ in the upper left corner of the page and choose **Management & Governance** > **Log Tank Service**.

Step 2 In the **Log Groups** area, locate a target log group and click its name.

Figure 16-9 Downloading logs

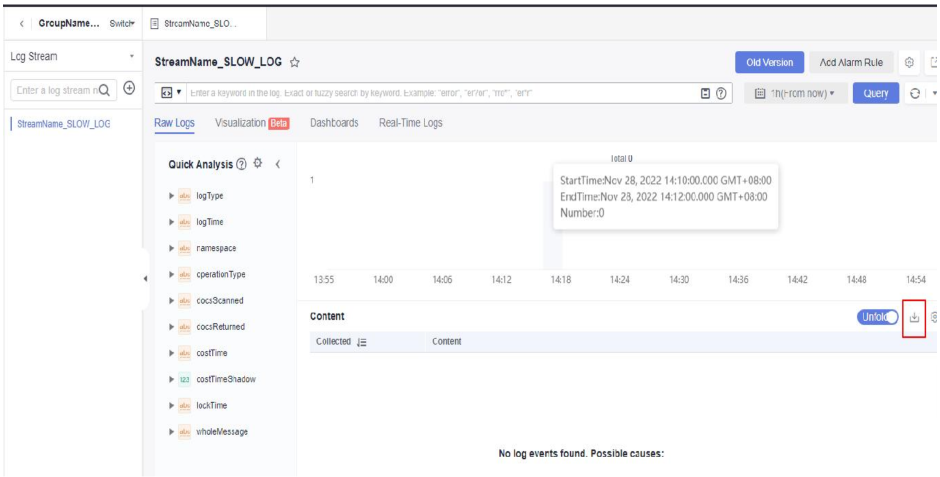

StartTime:Nov 28, 2022 14:10:00.000 GMT+08:00
EndTime:Nov 28, 2022 14:12:00.000 GMT+08:00
Number:0

No log events found. Possible causes:

Step 3 Click ⤓.

----End

16.2.2 Viewing Error Logs on the DDS Console

DDS log management allows you to view database-level logs, including warning- and error-level logs generated during database running, which help you analyze system problems.

Viewing and Exporting Log Details

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Error Logs**.

Step 6 On the displayed page, click **Error Logs**. Then, view the log details on the **Log Details** tab.

- For a cluster instance, you can view error logs of the dds mongos, shard, and config nodes.

Figure 16-10 Viewing error logs of a cluster instance

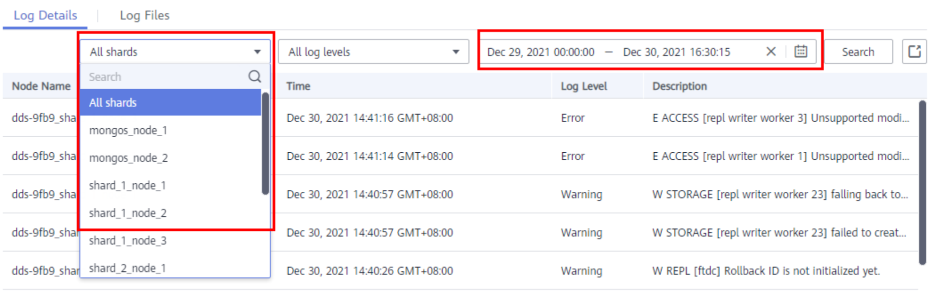

Node Name	Time	Log Level	Description
dds-9fb9_sha	Dec 30, 2021 14:41:16 GMT+08:00	Error	E ACCESS [repl writer worker 3] Unsupported modi...
dds-9fb9_sha	Dec 30, 2021 14:41:14 GMT+08:00	Error	E ACCESS [repl writer worker 1] Unsupported modi...
dds-9fb9_sha	Dec 30, 2021 14:40:57 GMT+08:00	Warning	W STORAGE [repl writer worker 23] falling back to...
dds-9fb9_shar	Dec 30, 2021 14:40:57 GMT+08:00	Warning	W STORAGE [repl writer worker 23] failed to creat...
dds-9fb9_shar	Dec 30, 2021 14:40:26 GMT+08:00	Warning	W REPL [ftdc] Rollback ID is not initialized yet.

- For a replica set instance, you can view the error logs of the primary, secondary, hidden nodes, and read replicas.

Figure 16-11 Viewing error logs of a replica set instance

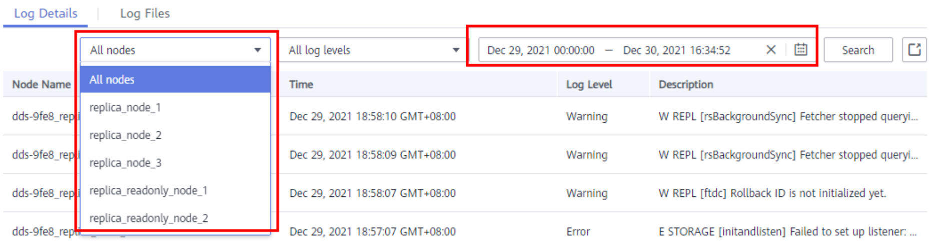

Node Name	Time	Log Level	Description
dds-9fe8_repl	Dec 29, 2021 18:58:10 GMT+08:00	Warning	W REPL [rsBackgroundSync] Fetcher stopped queryi...
dds-9fe8_repl	Dec 29, 2021 18:58:09 GMT+08:00	Warning	W REPL [rsBackgroundSync] Fetcher stopped queryi...
dds-9fe8_repl	Dec 29, 2021 18:58:07 GMT+08:00	Warning	W REPL [ftdc] Rollback ID is not initialized yet.
dds-9fe8_repl	Dec 29, 2021 18:57:07 GMT+08:00	Error	E STORAGE [initandlisten] Failed to set up listener: ...

- For a single node instance, you can view error logs of the current node.

Figure 16-12 Viewing error logs of a single-node instance

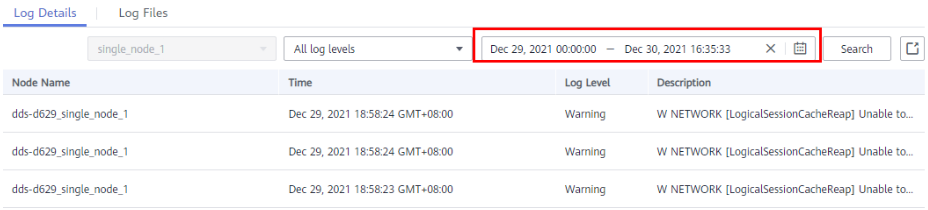

Log Details | Log Files

single_node_1 | All log levels | Dec 29, 2021 00:00:00 – Dec 30, 2021 16:35:33 | Search

Node Name	Time	Log Level	Description
dds-d629_single_node_1	Dec 29, 2021 18:58:24 GMT+08:00	Warning	W NETWORK [LogicalSessionCacheReap] Unable to...
dds-d629_single_node_1	Dec 29, 2021 18:58:24 GMT+08:00	Warning	W NETWORK [LogicalSessionCacheReap] Unable to...
dds-d629_single_node_1	Dec 29, 2021 18:58:23 GMT+08:00	Warning	W NETWORK [LogicalSessionCacheReap] Unable to...

- You can view up to 2,000 error logs of a specified node type, at a specified level, and within a specified period.

Step 7 On the **Log Details** tab, click **Advanced Search**.

Figure 16-13 Advanced search

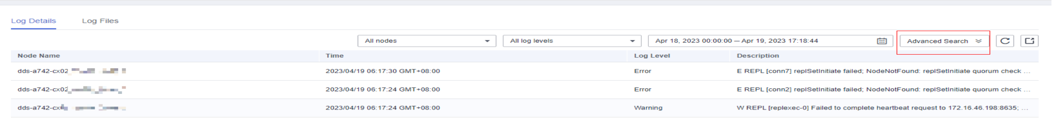

Log Details | Log Files

All nodes | All log levels | Apr 18, 2023 00:00:00 — Apr 19, 2023 17:18:44 | Advanced Search

Node Name	Time	Log Level	Description
dds-a742-cx02[redacted]	2023/04/19 06:17:30 GMT+08:00	Error	E REPL [conn7] replSetInitiate failed; NodeNotFound: replSetInitiate quorum check ...
dds-a742-cx02[redacted]	2023/04/19 06:17:24 GMT+08:00	Error	E REPL [conn2] replSetInitiate failed; NodeNotFound: replSetInitiate quorum check ...
dds-a742-cx[redacted]	2023/04/19 06:17:24 GMT+08:00	Warning	W REPL [replexec-0] Failed to complete heartbeat request to 172.16.46.198:8635; ...

Step 8 Specify **Keyword** and click **Search** to view log information.

Figure 16-14 Setting advanced search parameters

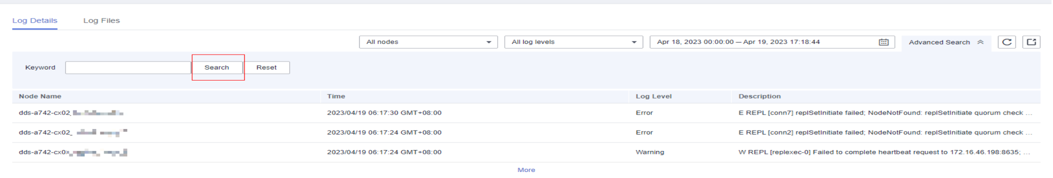

Log Details | Log Files

All nodes | All log levels | Apr 18, 2023 00:00:00 — Apr 19, 2023 17:18:44 | Advanced Search

Keyword | Search | Reset

Node Name	Time	Log Level	Description
dds-a742-cx02[redacted]	2023/04/19 06:17:30 GMT+08:00	Error	E REPL [conn7] replSetInitiate failed; NodeNotFound: replSetInitiate quorum check ...
dds-a742-cx02[redacted]	2023/04/19 06:17:24 GMT+08:00	Error	E REPL [conn2] replSetInitiate failed; NodeNotFound: replSetInitiate quorum check ...
dds-a742-cx0[redacted]	2023/04/19 06:17:24 GMT+08:00	Warning	W REPL [replexec-0] Failed to complete heartbeat request to 172.16.46.198:8635; ...

More

Step 9 To clear the parameter settings of **Advanced Search**, click **Reset**.

Figure 16-15 Resetting advanced search parameters

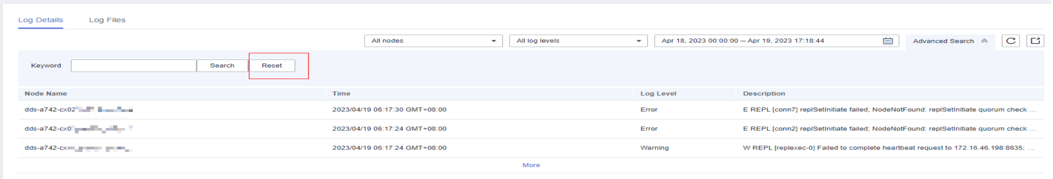

Log Details | Log Files

All nodes | All log levels | Apr 18, 2023 00:00:00 — Apr 19, 2023 17:18:44 | Advanced Search

Keyword | Search | Reset

Node Name	Time	Log Level	Description
dds-a742-cx02[redacted]	2023/04/19 06:17:30 GMT+08:00	Error	E REPL [conn7] replSetInitiate failed; NodeNotFound: replSetInitiate quorum check ...
dds-a742-cx0[redacted]	2023/04/19 06:17:24 GMT+08:00	Error	E REPL [conn2] replSetInitiate failed; NodeNotFound: replSetInitiate quorum check ...
dds-a742-cx[redacted]	2023/04/19 06:17:24 GMT+08:00	Warning	W REPL [replexec-0] Failed to complete heartbeat request to 172.16.46.198:8635; ...

More

Step 10 On the **Log Details** tab, click ↗ in the upper right corner of the log list to export log details.

- View the .csv file exported to your local PC.
- Up to 2,000 log details can be exported at a time.

----End

Downloading Logs

Step 1 Log in to the management console.

Step 2 Click ⊙ in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the Community Edition instance name.

Step 5 In the navigation pane on the left, choose **Error Logs**.

Step 6 On the **Error Logs** page, click the **Log Files** tab. Locate a log whose status is **Preparation completed** and click **Download** in the **Operation** column.

Figure 16-16 Error Logs

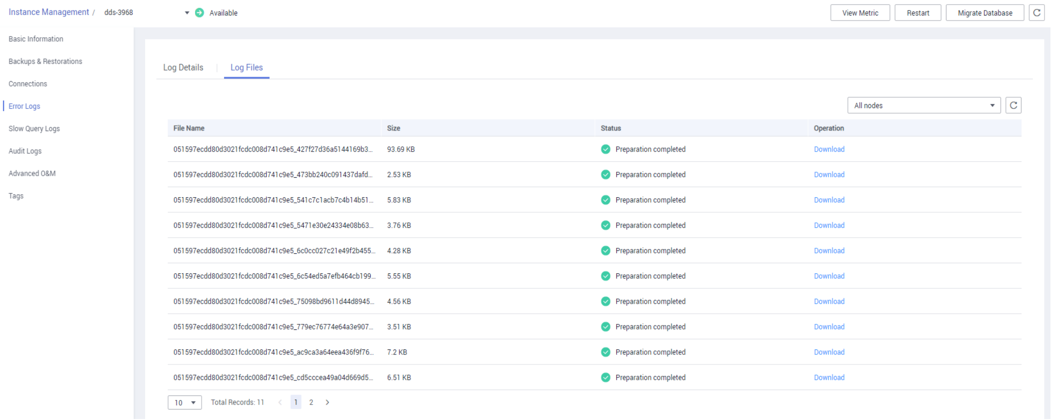

File Name	Size	Status	Operation
051597ecdd80d3021fcdc008d741c9e5_427f27d36a5144169b3...	93.69 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_473bb240c091437dafd...	2.53 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_541c7c1acb7c4b14b51...	5.83 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_5471e30e24334e08b63...	3.76 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_6c0cc027c21e49f2b455...	4.28 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_6c54ed5a7efb464cb199...	5.55 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_75098bd9611d44d8945...	4.56 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_779ec76774e64a3e907...	3.51 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_ac9ca3a64eea436f9f76...	7.2 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_cd5cccea49a04d669d5...	6.51 KB	Preparation completed	Download

10 Total Records: 11 1 2

- The system automatically loads the downloading preparation tasks. The time it takes to download the logs depends on the file size and on the network environment.
  - During the downloading preparation, the log status is **Preparing**.
  - Once the logs are ready for download, the log status changes to **Preparation completed**.
  - If the downloading preparation fails, the log status is **Abnormal**.
- You can download only one log file from a node. The maximum size of a log file to be downloaded is 40 MB.
- The download link is valid for 15 minutes. After the download link expires, a message is displayed indicating that the download link has expired. To download the log, click **OK**.

----End

16.3 Slow Query Logs

16.3.1 Viewing Slow Query Logs on the LTS Console

You can analyze, search for, monitor, download, and view real-time logs on the LTS console.

Querying Slow Query Logs Reported to LTS

NOTE

You have enabled log reporting to LTS. For details, see **Log Reporting**.

Step 1 Click ☰ in the upper left corner of the page and choose **Management & Governance** > **Log Tank Service**.

Step 2 In the **Log Groups** area, locate a target log group and click its name. For details about logs, see **Log Management**.

Figure 16-17 Viewing log details

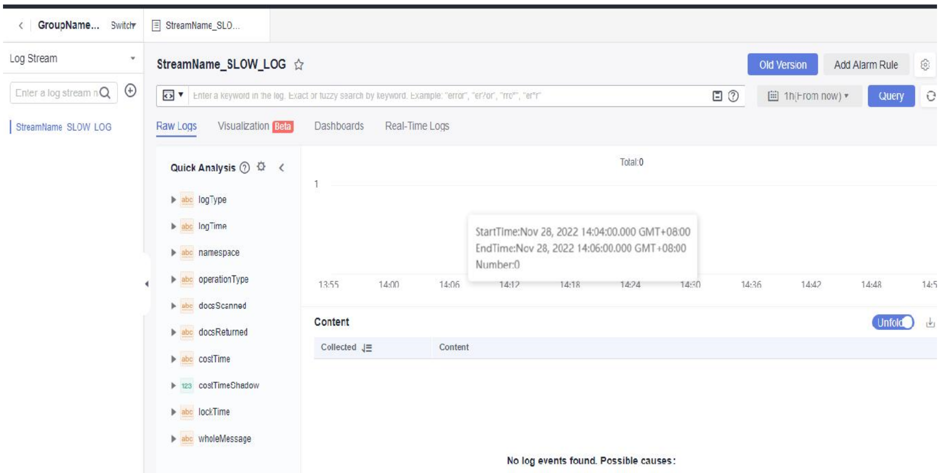

StartTime:Nov 28, 2022 14:04:00.000 GMT+08:00
EndTime:Nov 28, 2022 14:06:00.000 GMT+08:00
Number:0

No log events found. Possible causes:

----End

Downloading Slow Query Logs Reported to LTS

NOTE

If you have enabled log reporting to LTS for your DB instance in **Log Reporting**, you can download logs on the LTS console.

Step 1 Click ☰ in the upper left corner of the page and choose **Management & Governance** > **Log Tank Service**.

Step 2 In the **Log Groups** area, locate a target log group and click its name.

Figure 16-18 Downloading logs

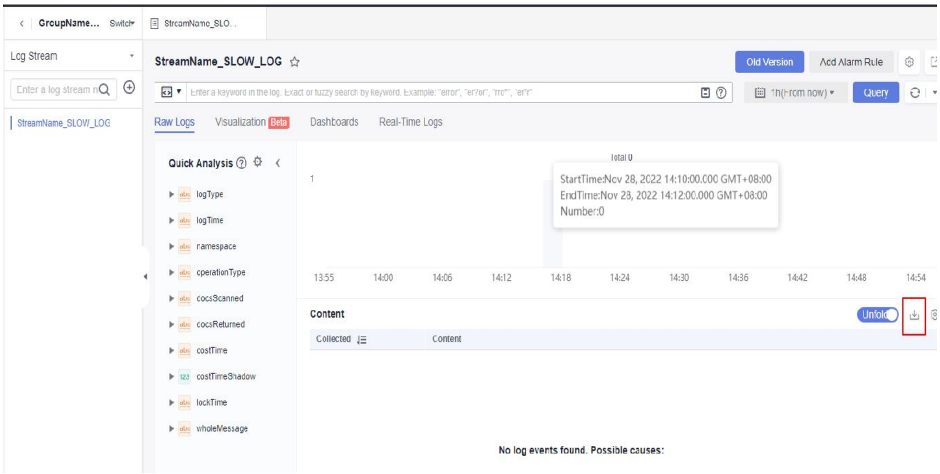

StartTime:Nov 28, 2022 14:10:00.000 GMT+08:00
EndTime:Nov 28, 2022 14:12:00.000 GMT+08:00
Number:0

No log events found. Possible causes:

Step 3 Click ⤓.

----End

16.3.2 Viewing Slow Query Logs on the DDS Console

Slow query logs record statements that exceed **operationProfiling.slowOpThresholdMs** (500 seconds by default). You can view log details and statistics to identify statements that are executing slowly and optimize the statements. You can also download slow query logs for service analysis.

Precautions

- Community Edition instances allow you to view and export log details, enable Show Original Log, and download log files on the management console.
- The Show Original Log function cannot be enabled when you delete DB instances, add nodes, change DB instance class, rebuild secondary node, or the DB instance is frozen.
- If **Show Original Log** is being enabled, you cannot delete instances, add nodes, or change instance class.
- By default, if the execution time of a SQL statement exceeds 500 ms, a slow query log is recorded.
- When the size of slow query logs reaches a specified threshold, old data is automatically deleted. If you need to analyze slow query logs, download the logs on the console in a timely manner.
- You can query slow logs for the last 30 days.
- You cannot delete slow query logs of DDS.
- When you export data on the **Log Details** page, all logs displayed on the current page will be exported.
- For details about how to sort slow query logs by field, such as execution completion time, SQL statement, client IP address, user, execution duration, lock wait time, scanned documents, returned documents, and scanned indexes, see "Slow Query Logs" in *Data Admin Service (DAS) User Guide*.

- Slow query logs may have a delay of several seconds to minutes, depending on the number of generated slow query logs and the DB instance load.
- In slow query log monitoring, the data of each monitored node is generated based on the total number of slow query logs generated 5 minutes before the time point.
- To apply for the advanced search permission, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console. DB instances of 230830 and later versions support the advanced search function.
- To apply for the permission to load 500 slow query logs at a time, submit an application by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console. If the node type is set to **All nodes** or **All shards**, you are advised to set the query time range to less than 10 minutes.
- Slow query logs do not have strict consistency. That is, the slow query logs displayed on the page may not include the full slow query logs. The system collects slow query logs periodically. If slow query logs are generated too frequently, the collection period may not cover all slow query logs.

Parameter description

Table 16-1 Parameters related to DDS slow query logs

Parameter	Description
operationProfiling.slowOpThresholdMs	Queries that exceed the threshold in the unit of ms are deemed slow. The default value is **500 ms**. Unless otherwise specified, keeping the default value is recommended.

Enabling Show Original Log

NOTE

- If **Show Original Log** is enabled, original logs are displayed. By default, the system automatically deletes original logs after 30 days, and the period cannot be changed.
- If the instance a slow query log belongs to is deleted, related logs are deleted along with it.
- **Show Original Log** can be disabled after it is enabled. The slow query logs reported before the function is disabled are displayed. The slow query logs reported after the function is disabled are not displayed.

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Slow Query Logs**.

Step 6 On the displayed page, click **Slow Query Logs**. Then, click ◯ on the **Log Details** tab.

Figure 16-19 Enabling Show Original Log

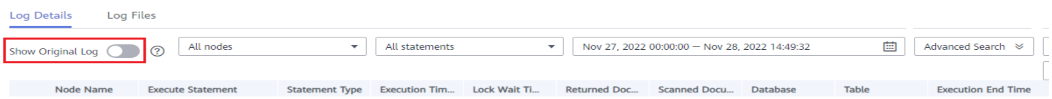

Step 7 In the displayed dialog box, click **Yes** to enable the function of slowing original logs.

----End

Viewing and Exporting Log Details

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Slow Query Logs**.

Step 6 On the **Slow Query Logs** page, set search criteria on the **Log Details** tab to view log information.

Figure 16-20 Querying slow query logs

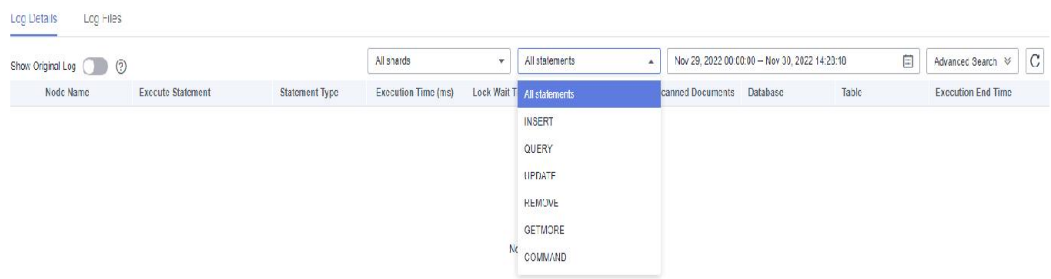

- Log records of all shards of a cluster instance
- Log records of all nodes in a replica set instance
- Slow query logs of a node in different time periods
- Slow query statements of the following levels
  - All statement type
  - INSERT

- QUERY
- UPDATE
- REMOVE
- GETMORE
- COMMAND

- You can view up to 2,000 slow logs of a specified node type, at a specified level, and within a specified period.

Step 7 On the **Log Details** tab, click **Advanced Search**.

Figure 16-21 Advanced search

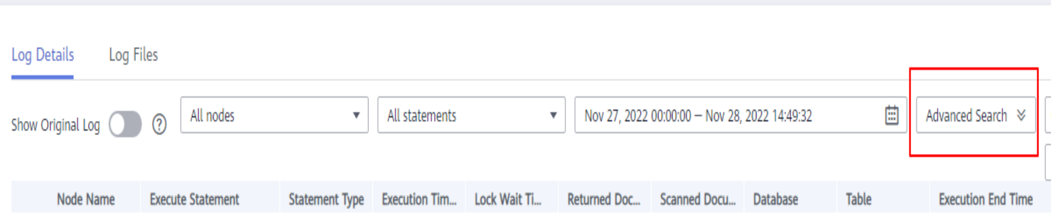

NOTE

- To apply for the advanced search permission, submit a service ticket by choosing **Service Tickets > Create Service Ticket** in the upper right corner of the management console. DB instances of 230830 and later versions support the advanced search function.

Step 8 Specify **Keyword**, **Execution Time (ms)**, **Returned Documents**, **Scanned Documents**, **Database**, **Table**, **Scaned Indexes**, **Username**, and **Client IP Address** and click **Search** to view log information.

Figure 16-22 Setting advanced search parameters

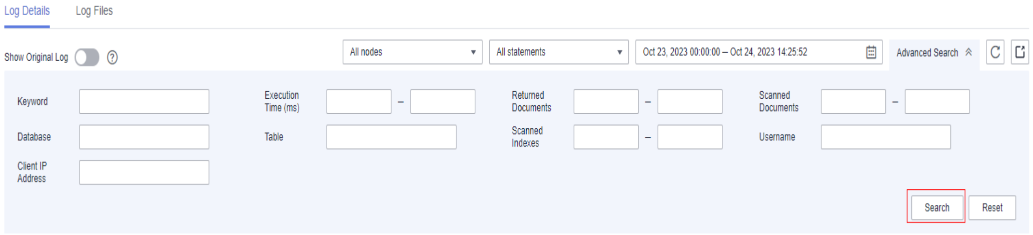

Step 9 To clear the parameter settings of **Advanced Search**, click **Reset**.

Figure 16-23 Resetting advanced search parameters

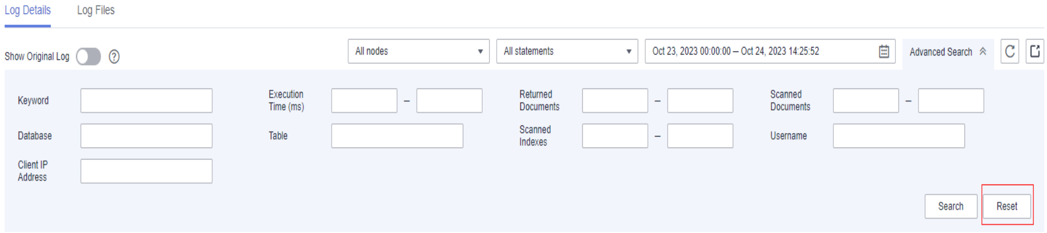

Step 10 On the **Log Details** tab, click ↗ in the upper right corner of the log list to export log details.

- View the .csv file exported to your local PC.
- Up to 2,000 log details can be exported at a time.

----End

Downloading Logs

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Slow Query Logs**.

Step 6 On the **Slow Query Logs** page, click the **Log Files** tab. Locate a log whose status is **Preparation completed** and click **Download** in the **Operation** column.

Figure 16-24 Slow Query Logs

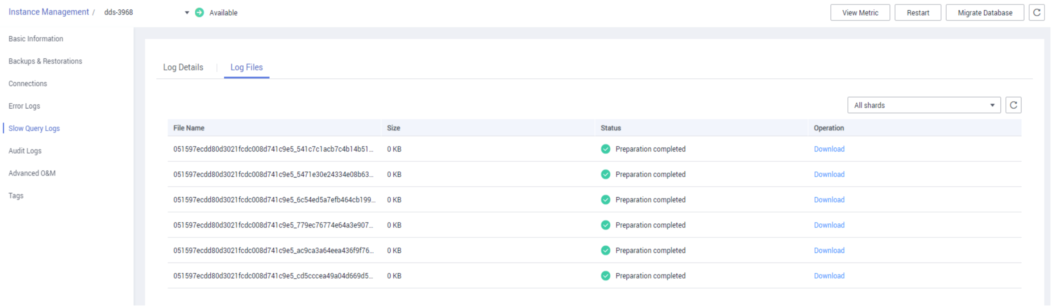

File Name	Size	Status	Operation
051597ecdd80d3021fcdc008d741c9e5_541c7c1acb7c4b14b51...	0 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_5471e30e24334e08b63...	0 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_6c54ed5a7efb464cb199...	0 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_779ec76774e64a3e907...	0 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_ac9ca3a64eea436f9f76...	0 KB	Preparation completed	Download
051597ecdd80d3021fcdc008d741c9e5_cd5cccea49a04d669d5...	0 KB	Preparation completed	Download

- The system automatically loads the downloading preparation tasks. The time required depends on the log file size and the network environment.
  - During the downloading preparation, the log status is **Preparing**.
  - Once the logs are ready for download, the log status changes to **Preparation completed**.
  - If the downloading preparation fails, the log status is **Abnormal**.
- You can download only one log file from a node. The maximum size of a log file to be downloaded is 40 MB.
- The download link is valid for 5 minutes. After the download link expires, a message is displayed indicating that the download link has expired. To download the log, click **OK**.

----End

Reference

How Do I Optimize Slow Operations?

16.4 Audit Logs

16.4.1 Audit Log Policy Management

An audit log records operations performed on your databases and collections. The generated log files are stored in OBS. Auditing logs can enhance your database security and help you analyze the cause of failed operations.

Precautions

- The audit policy of a DDS DB instance is disabled by default. You can enable it based on your service requirements. After the function is enabled, the system records audit information about read and write operations, which may deteriorate the performance by 15% to 20%.
- You will be charged for enabling SQL audit log. For details, see **Service Pricing**.
- DDS checks generated audit logs. If the retention period of logs exceeds the period you set, DDS will delete the logs. It is recommended that audit logs be stored for more than 180 days for tracing and problem analysis.
- After the audit policy is modified, DDS audits logs according to the new policy and the retention period of the original audit logs is subject to the modified retention period.
- You are not advised to delete audit logs. To delete audit logs, ensure that this operation meets external and internal security compliance requirements, and **download audit logs** and back them up locally. Audit logs cannot be restored after being deleted. Exercise caution when performing this operation.
- You can view, download, and delete DDS instance audit logs on the DDS console. For details, see **Viewing Audit Logs on the DDS Console**. By enabling log reporting in **Log Reporting**, you can also view details about audit logs of DDS DB instances on the LTS console, including searching for logs, monitoring logs, downloading logs, and viewing real-time logs. For details, see **Viewing Audit Logs on the LTS Console**.
- By default, audit logs are generated every hour. If the size of an audit log exceeds 10 MB, a new audit log is generated.
- Your data must be encoded in UTF-8 format. For data in other format, the auditing result of the corresponding statement may be missing or contain garbled characters.
- Audit log files stored on OBS are invisible to you. They are only visible in the DDS backend management system.

Example Traces

The following is an example of querying the replica set status. For details about the fields, see **Trace Structure**.

```
{
  "atype": "replSetGetStatus",
  "ts": {
    "$date": "2022-06-29T07:23:29.077+0000"
  },
  "local": {
    "ip": "127.0.0.1",
    "port": 8636
  },
  "remote": {
    "ip": "127.0.0.1",
    "port": 50860
  },
  "users": [
    {
      "user": "rwuser",
      "db": "admin"
    }
  ],
  "roles": [
    {
      "role": "root",
      "db": "admin"
    }
  ],
  "param": {
    "command": "replSetGetStatus",
    "ns": "admin",
    "args": {
      "replSetGetStatus": 1,
      "forShell": 1,
      "$clusterTime": {
        "clusterTime": {
          "$timestamp": {
            "t": 1656487409,
            "i": 117
          }
        },
        "signature": {
          "hash": {
            "$binary": "PTJhGQ6cr8RyzuqbevXfG0xWj/c=",
            "$type": "00"
          },
          "keyId": {
            "$numberLong": "7102437926763495425"
          }
        }
      },
      "$db": "admin"
    }
  },
  "result": 0
}
```

Configuring the Audit Policy

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, choose **Audit Logs**.

Step 6 On the **Audit Logs** page, click **Set Audit Policy**.

Step 7 On the displayed page, click .

Step 8 Configure required parameters and click **OK** to enable the audit policy.

Figure 16-25 Enabling audit policy

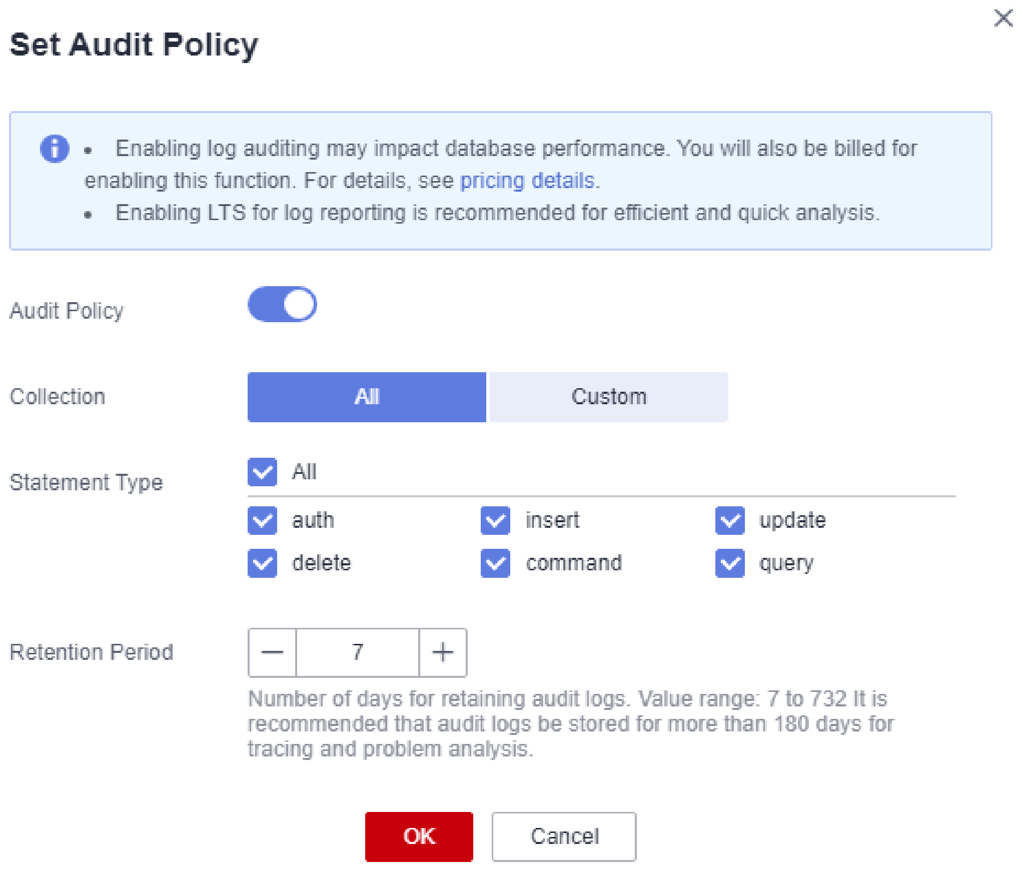

Set Audit Policy

- Enabling log auditing may impact database performance. You will also be billed for enabling this function. For details, see pricing details.
- Enabling LTS for log reporting is recommended for efficient and quick analysis.

Audit Policy ●

Collection: All | Custom

Statement Type: ☑ All

☑ auth ☑ insert ☑ update

☑ delete ☑ command ☑ query

Retention Period: − 7 +

Number of days for retaining audit logs. Value range: 7 to 732 It is recommended that audit logs be stored for more than 180 days for tracing and problem analysis.

Table 16-2 Parameter description

Parameter	Description
All	Audit all collections in the instance.

Parameter	Description
Custom	Audit specified databases or collections in the instance. The database or collection name cannot contain spaces or the following special characters: /\' : "[]{}() The dollar sign ($) can be used only as an escape character. The database name can contain a maximum of 64 characters. If you enter a combined database and collection name, the total name length is 120 characters with the database name length of no more than 64 characters and the collection name cannot be blank, contain **null**, or use **system.** in prefix.
Statement Type	You can query audit logs of specified statements in a collection, including auth, insert, update, delete, command and query statements.
Retention Days	The number of days to retain audit logs. Range: 7 to 732

- After the audit policy is enabled, you can modify it as required. After the modification, logs are generated according to the new policy and the retention period of the original logs is subject to the modified retention period.

  To modify the audit policy, click **Set Audit Policy**. In the dialog box that is displayed, modify the audit policy.

Figure 16-26 Modifying the audit policy

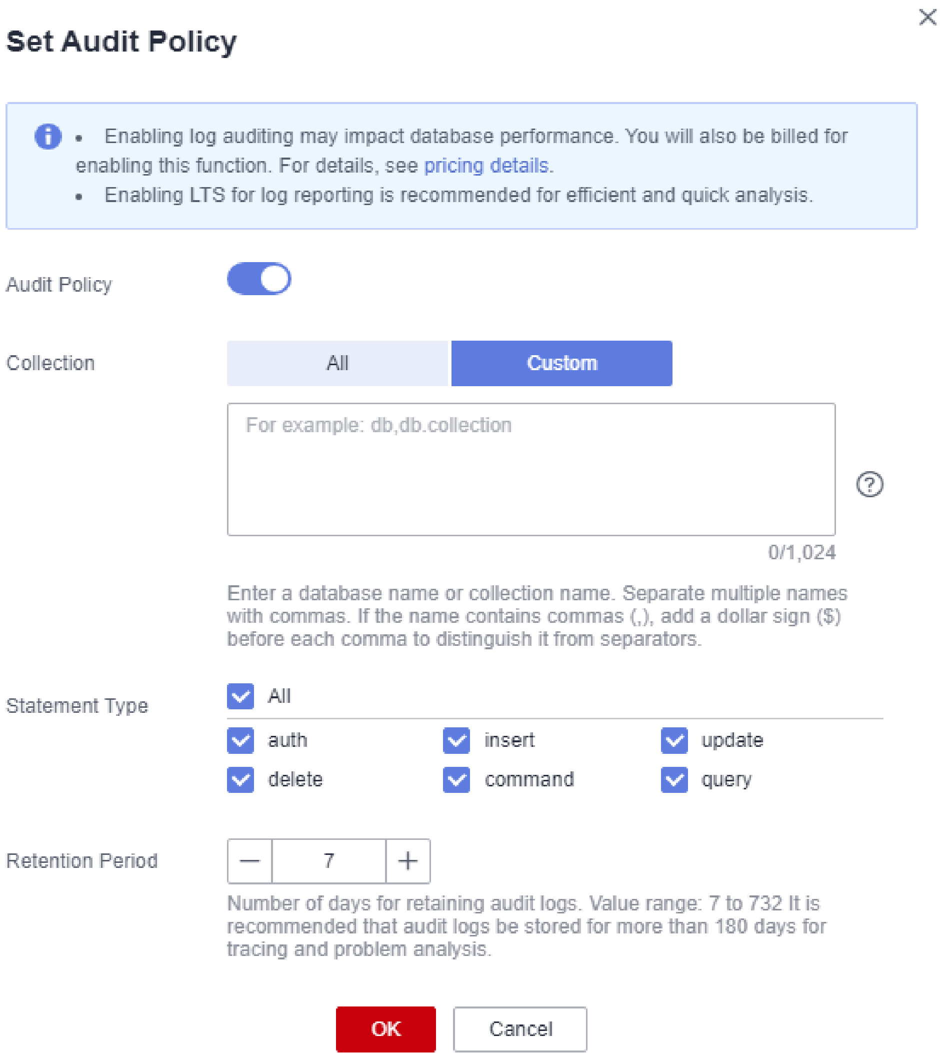

Set Audit Policy

- Enabling log auditing may impact database performance. You will also be billed for enabling this function. For details, see pricing details.
- Enabling LTS for log reporting is recommended for efficient and quick analysis.

Audit Policy: on

Collection: All | **Custom**

For example: db,db.collection

0/1,024

Enter a database name or collection name. Separate multiple names with commas. If the name contains commas (,), add a dollar sign ($) before each comma to distinguish it from separators.

Statement Type: ☑ All

☑ auth ☑ insert ☑ update

☑ delete ☑ command ☑ query

Retention Period: 7

Number of days for retaining audit logs. Value range: 7 to 732 It is recommended that audit logs be stored for more than 180 days for tracing and problem analysis.

- Disable the audit policy.

  NOTE

  After the audit policy is disabled, no audit log is generated.

  To disable the audit policy, click . **Figure 16-27** shows the dialog box for setting the backup policy.

Figure 16-27 Disabling audit policy

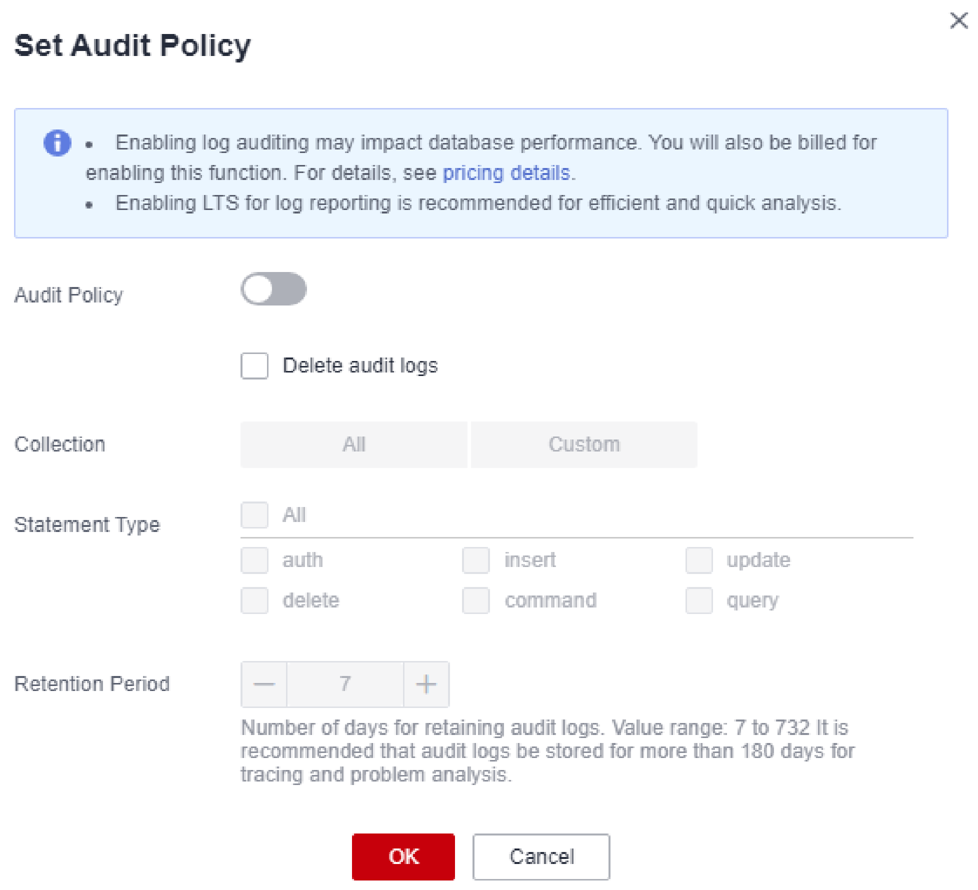

Set Audit Policy

- Enabling log auditing may impact database performance. You will also be billed for enabling this function. For details, see pricing details.
- Enabling LTS for log reporting is recommended for efficient and quick analysis.

Audit Policy

□ Delete audit logs

Collection: All | Custom

Statement Type: □ All

□ auth □ insert □ update

□ delete □ command □ query

Retention Period: — 7 +

Number of days for retaining audit logs. Value range: 7 to 732 It is recommended that audit logs be stored for more than 180 days for tracing and problem analysis.

You can determine whether to delete all audit logs:

– If you do not select **Delete audit logs**, all audit logs within the retention period will be retained. You can manually delete them later.

– If you select **Delete audit logs**, all audit logs within the retention period will be deleted.

Click **OK**.

----End

16.4.2 Viewing Audit Logs on the LTS Console

You can analyze, search for, monitor, download, and view real-time logs on the LTS console.

Querying Audit Logs Reported to LTS

NOTE

You have enabled log reporting to LTS. For details, see **Log Reporting**.

Step 1 Click ☰ in the upper left corner of the page and choose **Management & Governance** > **Log Tank Service**.

Step 2 In the **Log Groups** area, locate a target log group and click its name. For details about logs, see **Log Management**.

Figure 16-28 Viewing log details

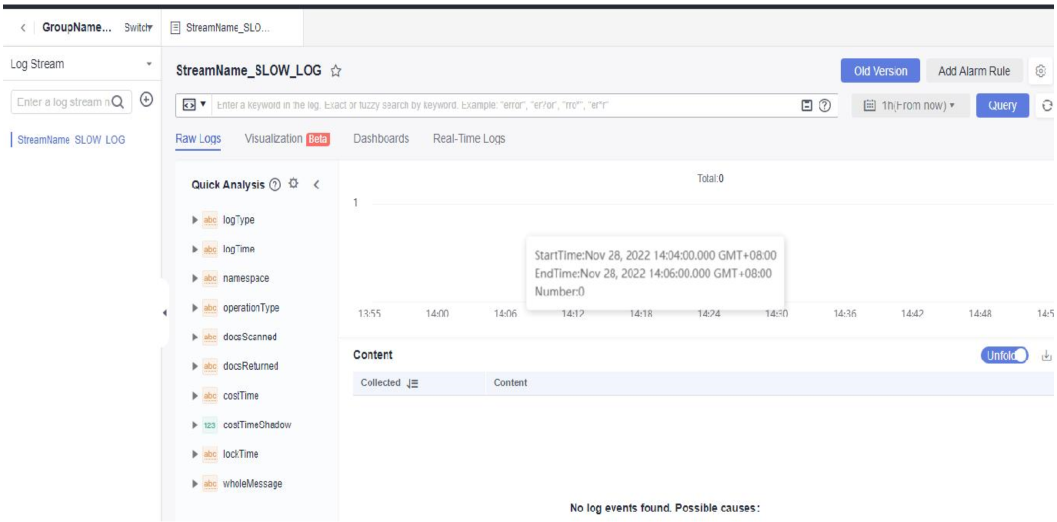

StartTime:Nov 28, 2022 14:04:00.000 GMT+08:00
EndTime:Nov 28, 2022 14:06:00.000 GMT+08:00
Number:0

No log events found. Possible causes:

----End

Downloading Audit Logs Reported to LTS

NOTE

If you have enabled log reporting to LTS for your DB instance in **Log Reporting**, you can download logs on the LTS console.

Step 1 Click ☰ in the upper left corner of the page and choose **Management & Governance** > **Log Tank Service**.

Step 2 In the **Log Groups** area, locate a target log group and click its name.

Figure 16-29 Downloading logs

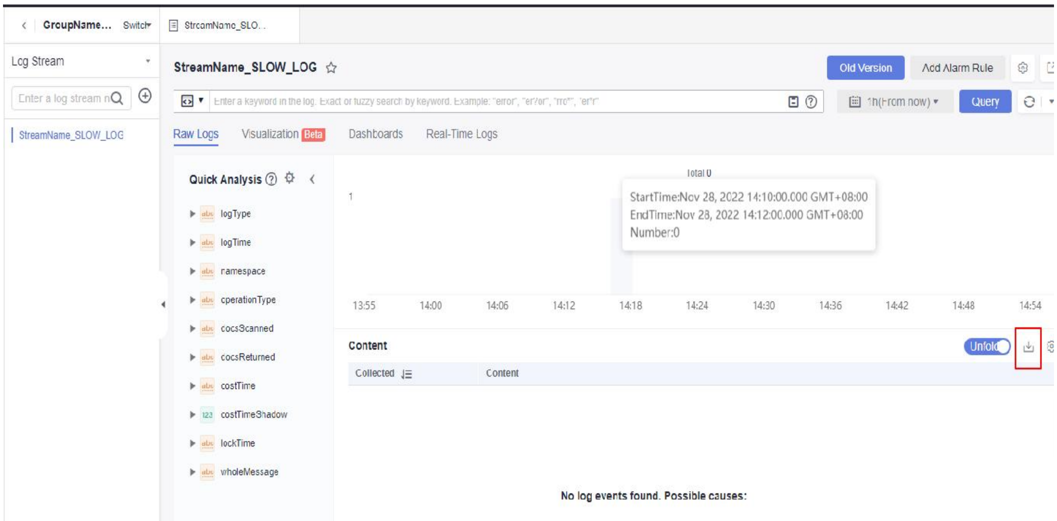

StartTime:Nov 28, 2022 14:10:00.000 GMT+08:00
EndTime:Nov 28, 2022 14:12:00.000 GMT+08:00
Number:0

No log events found. Possible causes:

Step 3 Click ⤓.

----End

16.4.3 Viewing Audit Logs on the DDS Console

You can view, download, and delete audit logs on the DDS console.

Querying Audit Logs

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate a target DB instance and click its name.

Step 5 In the navigation pane on the left, choose **Audit Logs**.

Step 6 On the **Audit Logs** page, locate a target log file and click **Download** in the **Operation** column to download the log file to the local PC for query.

----End

Downloading Logs

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate a target DB instance and click its name.

Step 5 In the navigation pane on the left, choose **Audit Logs**.

Step 6 On the **Audit Logs** page, locate a target log file and click **Download** in the **Operation** column.

- The system automatically loads the downloading preparation tasks. The time required depends on the log file size and the network environment.
- The download link is valid for 5 minutes. After the download link expires, a message is displayed indicating that the download link has expired. To download the log, click **OK**.

----End

Deleting Logs

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate a target DB instance and click its name.

Step 5 In the navigation pane on the left, choose **Audit Logs**.

Step 6 On the **Audit Logs** page, locate a target log file and click **Delete** in the **Operation** column.

Step 7 Click **Yes**.

----End

17 Task Center

This section describes how to view the progress and result of asynchronous tasks on the **Task Center** page.

Precautions

Tasks that fail to be executed will be retained for seven days by default.

Tasks Overview

Table 17-1 List of tasks that can be viewed

Task Name	Description
Creating an instance	Creating a cluster instance or replica set instance.
Scaling up storage space	Scaling up the storage space of the shard node of a cluster instance or the storage space of a replica set instance.
Changing instance class	Changing the class of a cluster instance or replica set instance.
Adding nodes	Adding nodes to a cluster instance.
Adding read replicas	Adding read replicas to a cluster or replica set instance of Community Edition.
Restarting DB instances	Restarting a cluster instance, one or more cluster instance nodes, or a replica set instance.
Restoring to a new DB instance	Restoring data to a new cluster instance or replica set instance.
Restoring data to the original DB instance	Restoring data to a new Community Edition cluster instance, single node instance, or replica set instance.

Task Name	Description
Restoring to a point in time	Restoring a replica set instance to a point in time.
Restoring databases and tables to a point in time	Restores table-level data of a replica set instance to a specified point in time.
Performing a primary/ standby switchover	Perform a primary/standby switchover for a replica set instance.
Binding and unbinding an EIP	Bind or unbind an EIP to or from a cluster instance, single node instance, or replica set instance.
Switching SSL	Enabling or disabling SSL for a cluster instance, single node instance, or replica set instance.
Changing a database port	Changing the database port of a cluster , single node, or replica set instance of Community Edition.
Changing a security group	Changing the security group of a cluster, single node, or replica set instance of Community Edition.
Changing a private IP address	Changing the private IP address of a cluster, single node, or replica set instance of Community Edition.
Changing an AZ	Changing the AZ of a cluster, single node, or replica set instance of Community Edition.
Enabling the shard/config IP address	Enabling the shard/config address for the cluster instance of Community Edition.
Modifying the oplog size	Changing the oplog size of a cluster, single node, or replica set instance of Community Edition
Physical backup	Creating automated and manual backups of a cluster, single node, or replica set instance of Community Edition
Upgrading minor version	Community Edition cluster and replica set instances are being patched.

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 In the navigation pane on the left, click **Task Center**.

Step 5 In the navigation pane on the left, choose **Task Center**. Then, view the task progresses and results.

- You can view tasks in a specified period.
- The tasks can be located by DB instance name and ID or by task status or type from the drop-down list in the upper right corner.

----End

18 DBA Assistant

18.1 Real-Time Diagnosis

18.1.1 Real-Time Sessions

You can manage sessions in the following scenarios:

- Emergency Channel: If the maximum number of connections for an instance has been reached and the instance cannot be logged in to, you can view and kill unnecessary sessions through the emergency channel.
- History Logs: You can view history logs to learn details of the kill operations that you performed using the emergency channel function.

Precautions

- This function is not recommended unless you really need it. All your kill operations will be logged.
- DB instances of Community Edition 3.4, 4.0, 4.2, and 4.4 are supported.
- DB instances in the creating, frozen or abnormal state are not supported.
- Killing inactive sessions is not allowed.
- Real-time sessions are generated based on **currentOp** of a database at the current time point. If the execution time of a session is too short (less than or equal to milliseconds), you are not advised to view real-time sessions. If you need to collect statistics on all operations, see **Audit Log Policy Management**.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the cluster instance name.

Step 5 In the navigation tree, choose **DBA Assistant**.

Step 6 Choose **Real-Time Diagnosis**.

Step 7 Click **Real-Time Sessions**.

Step 8 On the displayed **Emergency Channel** page, view session statistics of the current instance node by overview, source, or namespace.

Step 9 By default, sessions are sorted and displayed in the session list in descending order by duration. You can also search for sessions by specifying **Sessions lasting longer than** or **Namespace**.

Step 10 Select the sessions that you want to kill and click **Kill Session**.

Figure 18-1 Killing a session

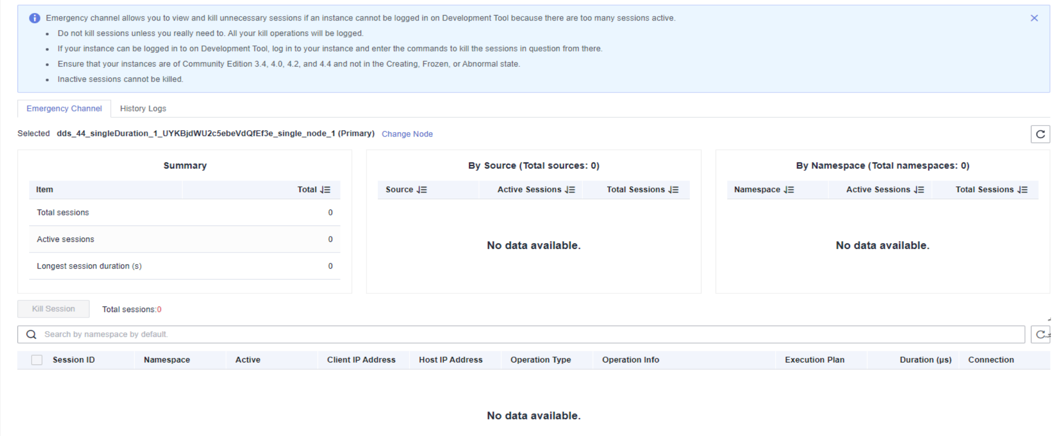

Emergency channel allows you to view and kill unnecessary sessions if an instance cannot be logged in on Development Tool because there are too many sessions active.

- Do not kill sessions unless you really need to. All your kill operations will be logged.
- If your instance can be logged in to on Development Tool, log in to your instance and enter the commands to kill the sessions in question from there.
- Ensure that your instances are of Community Edition 3.4, 4.0, 4.2, and 4.4 and not in the Creating, Frozen, or Abnormal state.
- Inactive sessions cannot be killed.

Selected dds_44_singleDuration_1_UYKBjdWU2c5ebeVdQfEf3e_single_node_1 (Primary) Change Node

Summary

Item	Total
Total sessions	0
Active sessions	0
Longest session duration (s)	0

By Source (Total sources: 0)

Source	Active Sessions	Total Sessions

No data available.

By Namespace (Total namespaces: 0)

Namespace	Active Sessions	Total Sessions

No data available.

Total sessions:0

Search by namespace by default.

□ Session ID	Namespace	Active	Client IP Address	Host IP Address	Operation Type	Operation Info	Execution Plan	Duration (μs)	Connection

No data available.

Step 11 In the **Kill Session** dialog box, confirm the session information and click **Yes**.

Step 12 Click **History Logs** to view the sessions killed through the emergency channel.

Figure 18-2 Viewing history logs

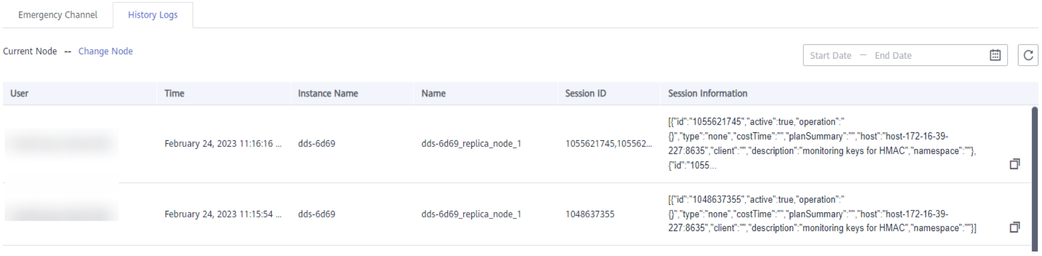

Current Node -- Change Node

Start Date – End Date

User	Time	Instance Name	Name	Session ID	Session Information
[redacted]	February 24, 2023 11:16:16 ...	dds-6d69	dds-6d69_replica_node_1	1055621745,105562...	[{"id":"1055621745","active":true,"operation":"{}","type":"none","costTime":"","planSummary":"","host":"host-172-16-39-227:8635","client":"","description":"monitoring keys for HMAC","namespace":""},{"id":"1055...
[redacted]	February 24, 2023 11:15:54 ...	dds-6d69	dds-6d69_replica_node_1	1048637355	[{"id":"1048637355","active":true,"operation":"{}","type":"none","costTime":"","planSummary":"","host":"host-172-16-39-227:8635","client":"","description":"monitoring keys for HMAC","namespace":""}]

----End

19 SQL Execution Control

Scenarios

- All requests whose execution duration exceeds *n* seconds need to be killed.
- Requests from an IP address for a specific client need to be killed.
- All requests for full table scan need to be killed.

Precautions

- The instance node must have 4 or more vCPUs.
- This function is available for replica set instances and cluster instances of version 3.4 or later.
- A maximum of 10 rules can be created for a DB instance.
- For an ultra-large cluster with more than 32 shards, creating and enabling rules whose **Node Type** is **shard** or **dds mongos_shard** will fail. You are advised to create rules whose **Node Type** is **dds mongos**.
- For a cluster with more than 10 shards, you are advised to select one rule at a time when enabling or disabling rules.
- This function is available only to whitelisted users. To use this function, you need to submit a service ticket. In the upper right corner of the management console, choose **Service Tickets > Create Service Ticket** to submit a service ticket.

Creating a Rule

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click its name.

Step 5 In the navigation tree on the left, click **SQL Execution Control**.

Figure 19-1 SQL Execution Control

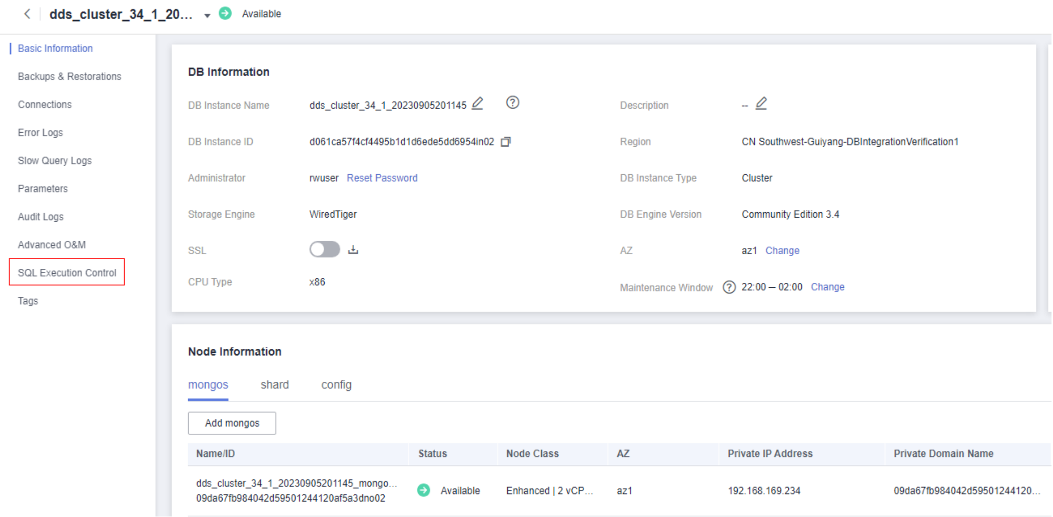

DB Information

DB Instance Name	dds_cluster_34_1_20230905201145	Description	--
DB Instance ID	d061ca57f4cf4495b1d1d6ede5dd6954in02	Region	CN Southwest-Guiyang-DBIntegrationVerification1
Administrator	rwuser Reset Password	DB Instance Type	Cluster
Storage Engine	WiredTiger	DB Engine Version	Community Edition 3.4
SSL		AZ	az1 Change
CPU Type	x86	Maintenance Window	22:00 – 02:00 Change

Node Information

mongos shard config

Add mongos

Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name
dds_cluster_34_1_20230905201145_mongo... 09da67fb984042d59501244120af5a3dno02	Available	Enhanced | 2 vCP...	az1	192.168.169.234	09da67fb984042d59501244120...

Step 6 Click **Create Rule**.

Figure 19-2 Create Rule

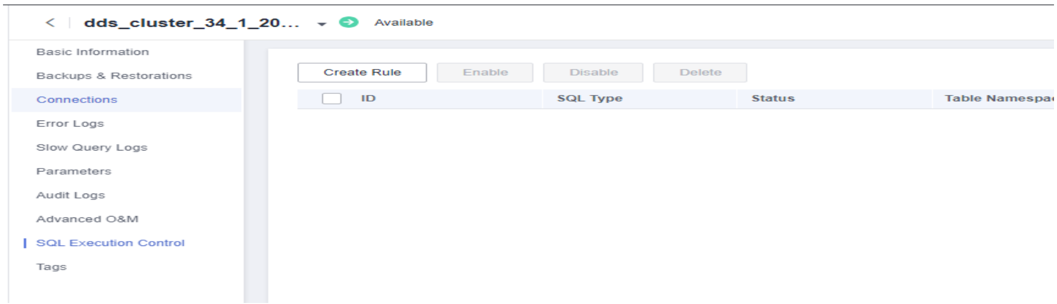

Step 7 On the **Create Rule** page, set parameters as required. For details, see **Table 19-1**.

Figure 19-3 Parameters for creating a rule

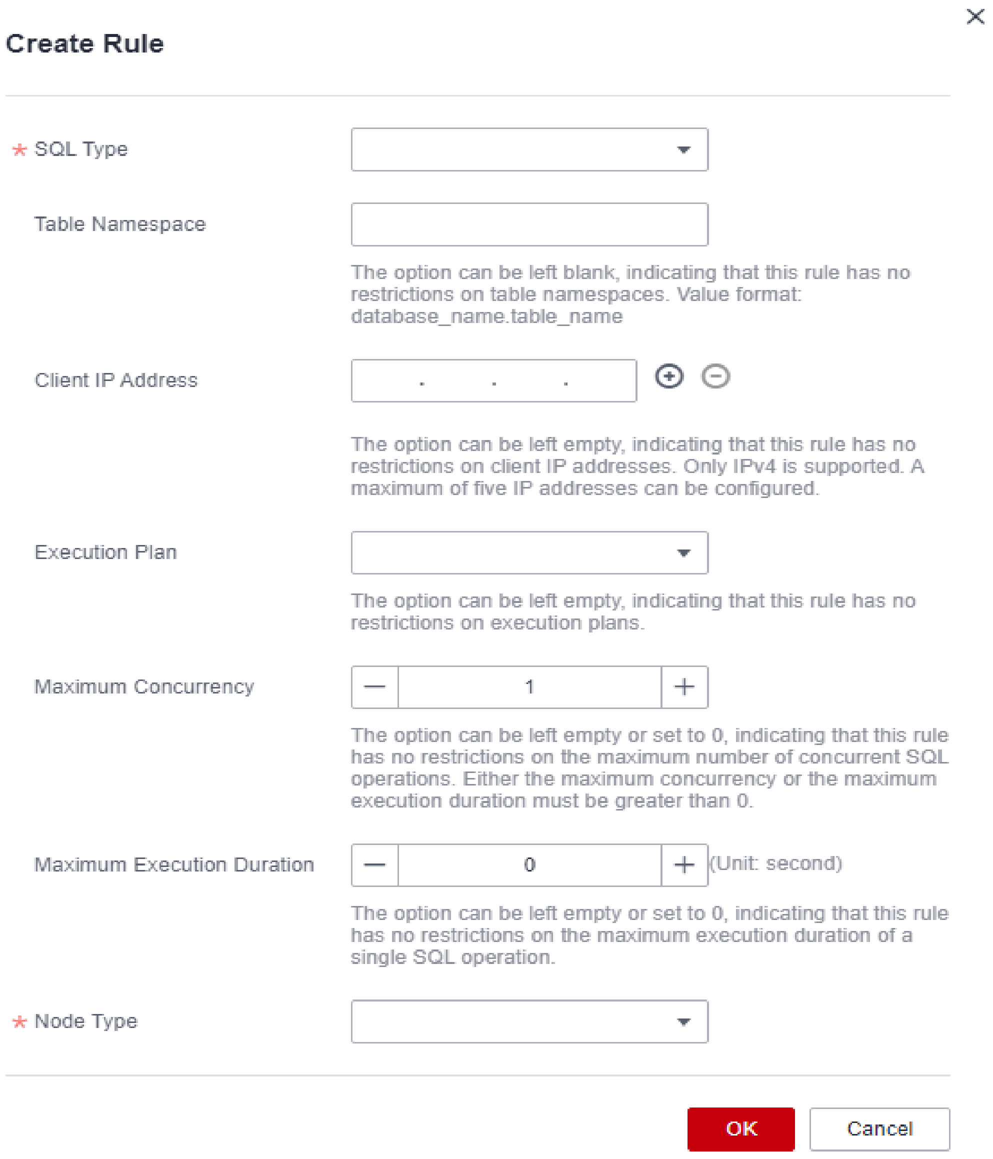

Create Rule

* SQL Type

Table Namespace

The option can be left blank, indicating that this rule has no restrictions on table namespaces. Value format: database_name.table_name

Client IP Address . . .

The option can be left empty, indicating that this rule has no restrictions on client IP addresses. Only IPv4 is supported. A maximum of five IP addresses can be configured.

Execution Plan

The option can be left empty, indicating that this rule has no restrictions on execution plans.

Maximum Concurrency 1

The option can be left empty or set to 0, indicating that this rule has no restrictions on the maximum number of concurrent SQL operations. Either the maximum concurrency or the maximum execution duration must be greater than 0.

Maximum Execution Duration 0 (Unit: second)

The option can be left empty or set to 0, indicating that this rule has no restrictions on the maximum execution duration of a single SQL operation.

* Node Type

Table 19-1 Parameter description

Parameter	Description
SQL Type	You can specify one or more SQL statement types. The value can be: ● **query**: operation for querying data. ● **update**: operation for updating data. ● **insert**: operation for inserting data. ● **remove**: operation for deleting data. ● **command**: command operation. ● **getmore**: operation for obtaining more data.

Parameter	Description
Table Namespace	● If this parameter is left blank, this rule applies to operations on all databases and tables in the DB instance. ● If this parameter is set to a database name, this rule applies to operations on all collections in the database. For example, the value can be **db1**. ● If this parameter is set to a value in the format of database_name.collection_name, this rule only applies to operations on the collection. For example, the value can be **db1.coll1**.
Client IP Address	If an ECS on Huawei Cloud is used, the value is the private IP address of the ECS. **NOTE** ● This parameter does not take effect for cluster DB instances of version 3.4.
Execution Plan	● By default, this parameter is left blank, indicating that this rule applies to all execution plans. ● The value **COLLSCAN** indicates that the operation for full table scan will be killed.
Maximum Concurrency	● The value **0** indicates that this parameter does not take effect. ● For example, if this parameter is set to **100**, a maximum of 100 operations that meet the conditions can be performed. **NOTE** If there are 110 **currentOp** operations that meet the conditions, 10 operations will be randomly killed. ● Either this parameter or **Maximum Execution Duration** must be greater than **0**.
Maximum Execution Duration	● The value **0** indicates that this parameter does not take effect. ● For example, if this parameter is set to **5**, **currentOp** operations executed for more than 5s will be killed. The value must be no less than **2**. ● Either this parameter or **Maximum Concurrency** must be greater than **0**.
Node Type	● **dds mongos** indicates that this rule only applies to mongos nodes in a DDS instance. ● **shard** indicates that this rule only applies to shard nodes. ● **dds mongos_shard** indicates that this rule applies to both mongos and shard nodes in a DDS instance. ● **replica** indicates that this rule applies to replica sets.

Step 8 Click **OK**.

----End

Enabling a Rule

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click its name.

Step 5 In the navigation tree on the left, click **SQL Execution Control**.

Figure 19-4 SQL Execution Control

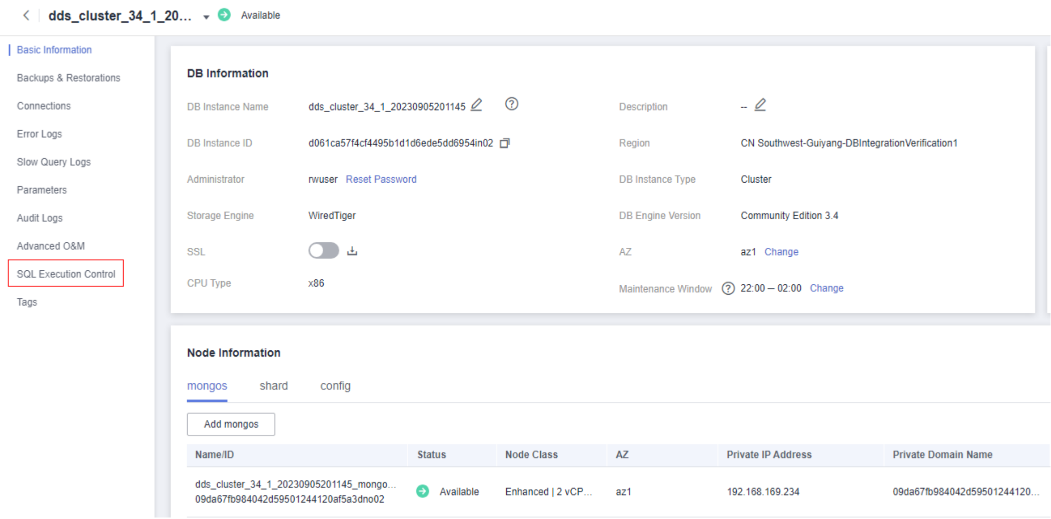

DB Information

DB Instance Name	dds_cluster_34_1_20230905201145	Description	--
DB Instance ID	d061ca57f4cf4495b1d1d6ede5dd6954in02	Region	CN Southwest-Guiyang-DBIntegrationVerification1
Administrator	rwuser Reset Password	DB Instance Type	Cluster
Storage Engine	WiredTiger	DB Engine Version	Community Edition 3.4
SSL		AZ	az1 Change
CPU Type	x86	Maintenance Window	22:00 – 02:00 Change

Node Information

Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name
dds_cluster_34_1_20230905201145_mongo... 09da67fb984042d59501244120af5a3dno02	Available	Enhanced | 2 vCP...	az1	192.168.169.234	09da67fb984042d59501244120...

Step 6 Locate the target rule and click **Enable** in the **Operation** column.

Figure 19-5 Enable

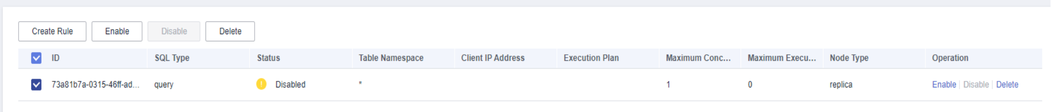

☑ ID	SQL Type	Status	Table Namespace	Client IP Address	Execution Plan	Maximum Conc...	Maximum Execu...	Node Type	Operation
☑ 73a81b7a-0315-46ff-ad...	query	Disabled	*			1	0	replica	Enable | Disable | Delete

Step 7 Click **Yes**.

Figure 19-6 Enable Rule

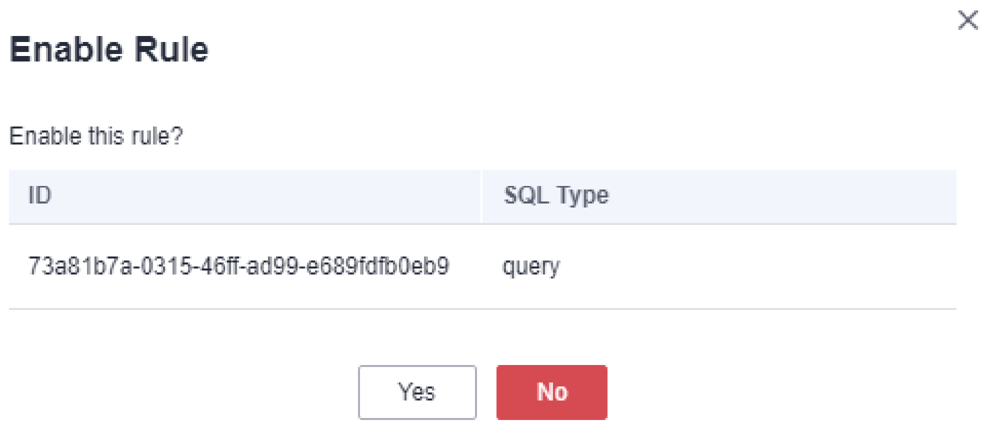

Enable Rule

Enable this rule?

ID	SQL Type
73a81b7a-0315-46ff-ad99-e689fdfb0eb9	query

Step 8 View the rule status on the **SQL Execution Control** page.

Figure 19-7 Status

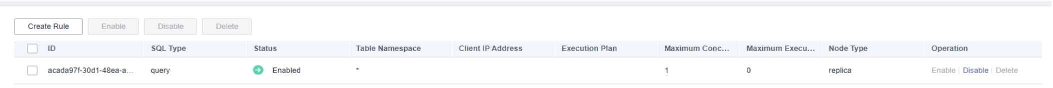

☐ ID	SQL Type	Status	Table Namespace	Client IP Address	Execution Plan	Maximum Conc...	Maximum Execu...	Node Type	Operation
☐ acada97f-30d1-48ea-a...	query	Enabled	*			1	0	replica	Enable | Disable | Delete

----End

Disabling a Rule

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click its name.

Step 5 In the navigation tree on the left, click **SQL Execution Control**.

Figure 19-8 SQL Execution Control

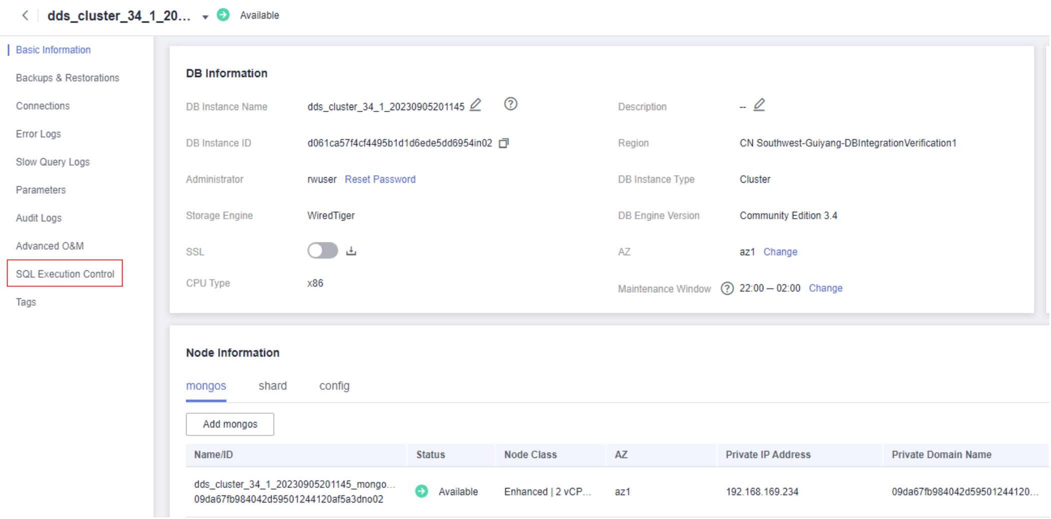

DB Information

DB Instance Name	dds_cluster_34_1_20230905201145	Description	--
DB Instance ID	d061ca57f4cf4495b1d1d6ede5dd6954in02	Region	CN Southwest-Guiyang-DBIntegrationVerification1
Administrator	rwuser	DB Instance Type	Cluster
Storage Engine	WiredTiger	DB Engine Version	Community Edition 3.4
SSL		AZ	az1
CPU Type	x86	Maintenance Window	22:00 – 02:00

Node Information

Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name
dds_cluster_34_1_20230905201145_mongo... 09da67fb984042d59501244120af5a3dno02	Available	Enhanced | 2 vCP...	az1	192.168.169.234	09da67fb984042d59501244120...

Step 6 Locate the target rule and click **Disable** in the **Operation** column.

Figure 19-9 Disable

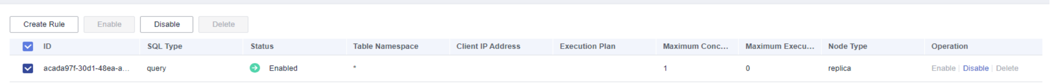

☑ ID	SQL Type	Status	Table Namespace	Client IP Address	Execution Plan	Maximum Conc...	Maximum Execu...	Node Type	Operation
☑ acada97f-30d1-48ea-a...	query	Enabled	*			1	0	replica	Enable | Disable | Delete

Step 7 Click **Yes**.

Figure 19-10 Disable Rule

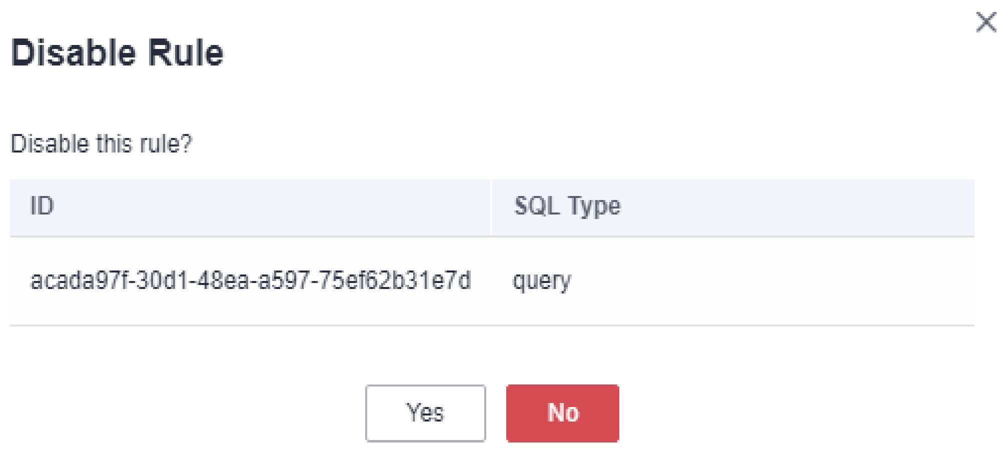

Disable Rule

Disable this rule?

ID	SQL Type
acada97f-30d1-48ea-a597-75ef62b31e7d	query

Step 8 View the rule status on the **SQL Execution Control** page.

Figure 19-11 Status

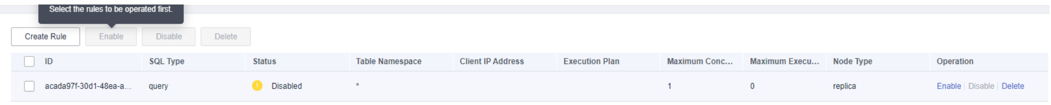

□ ID	SQL Type	Status	Table Namespace	Client IP Address	Execution Plan	Maximum Conc...	Maximum Execu...	Node Type	Operation
□ acada97f-30d1-48ea-a...	query	Disabled	*			1	0	replica	Enable | Disable | Delete

----End

Deleting a Rule

CAUTION

An enabled rule cannot be deleted. To delete a rule, you must disable the rule by referring to **Disabling a Rule**.

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

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, locate the target DB instance and click its name.

Step 5 In the navigation tree on the left, click **SQL Execution Control**.

Figure 19-12 SQL Execution Control

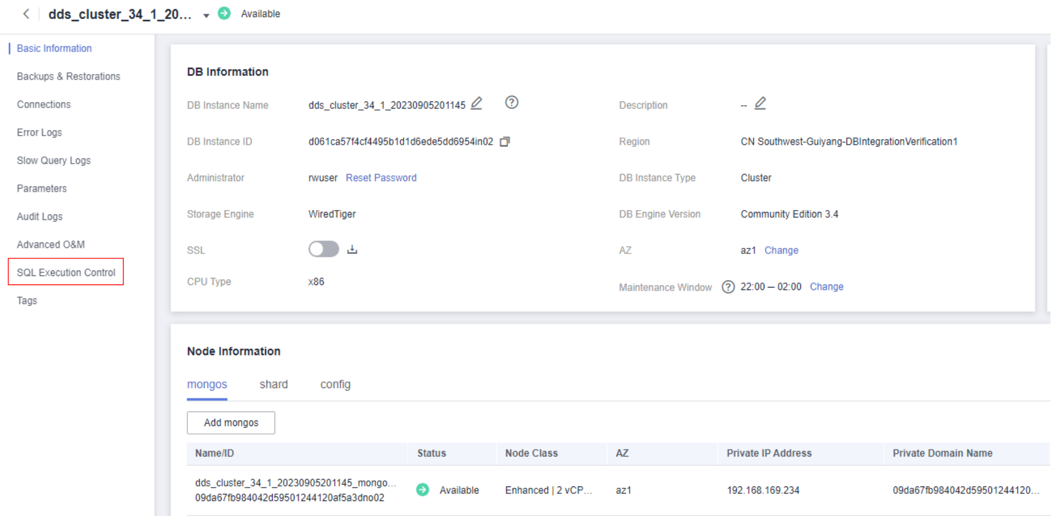

DB Information

DB Instance Name	dds_cluster_34_1_20230905201145	Description	--
DB Instance ID	d061ca57f4cf4495b1d1d6ede5dd6954in02	Region	CN Southwest-Guiyang-DBIntegrationVerification1
Administrator	rwuser Reset Password	DB Instance Type	Cluster
Storage Engine	WiredTiger	DB Engine Version	Community Edition 3.4
SSL		AZ	az1 Change
CPU Type	x86	Maintenance Window	22:00 – 02:00 Change

Node Information

Name/ID	Status	Node Class	AZ	Private IP Address	Private Domain Name
dds_cluster_34_1_20230905201145_mongo... 09da67fb984042d59501244120af5a3dno02	Available	Enhanced | 2 vCP...	az1	192.168.169.234	09da67fb984042d59501244120...

Step 6 Locate the target rule and click **Delete** in the **Operation** column.

Figure 19-13 Delete

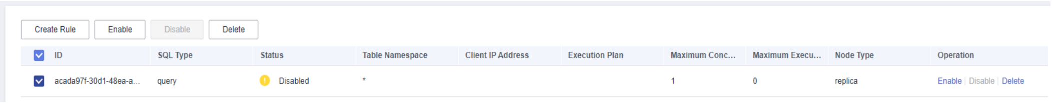

☑ ID	SQL Type	Status	Table Namespace	Client IP Address	Execution Plan	Maximum Conc...	Maximum Execu...	Node Type	Operation
☑ acada97f-30d1-48ea-a...	query	Disabled	*			1	0	replica	Enable | Disable | Delete

Step 7 Click **Yes**.

Figure 19-14 Delete Rule

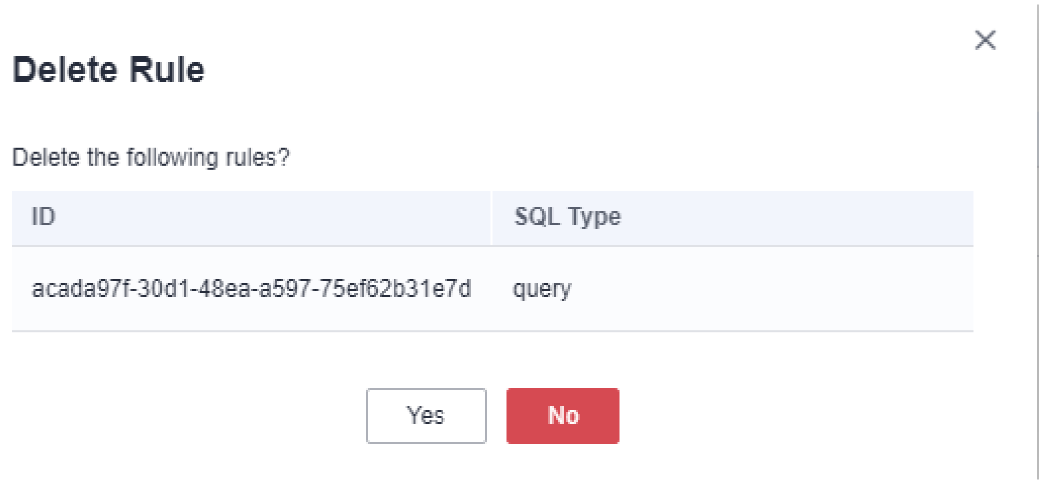

Delete Rule

Delete the following rules?

ID	SQL Type
acada97f-30d1-48ea-a597-75ef62b31e7d	query

----End

20 Cross-AZ Disaster Recovery

20.1 Creating a Cross-AZ Cluster Instance

DDS allows you to create a multi-AZ cluster. A multi-AZ cluster has higher DR capabilities than a single-AZ cluster and can withstand the impact caused by equipment room faults. To obtain higher DR capability, deploy resources across different AZs in the same region. If the AZ where the primary node is located fails due to power supply or network exceptions, the HA system automatically triggers a failover to ensure service continuity of a cluster instance.

This section describes how to create a multi-AZ cluster instance.

Precautions

- Only some regions support multi-AZ cluster instances.
- To create a multi-AZ instance, ensure that there are three or more AZs available in the region.
- Multi-AZ deployment means that the components of an instance are deployed in three different AZs.

Deployment Architecture Comparison

- Single AZ

  If an instance is deployed in a single AZ, all components of the instance are deployed in the same AZ. By default, anti-affinity deployment is configured. With an anti-affinity deployment, your primary, secondary, and hidden nodes are deployed on different physical machines for high availability.
- Multiple AZs

  The components of an instance are deployed in three different AZs for disaster recovery.
  - Two dds mongos nodes are respectively deployed in two AZs. If one dds mongos node is added, it will be deployed in the third AZ.
  - The primary, secondary, and hidden shard nodes are randomly and evenly deployed in three AZs.

Figure 20-1 Multi-AZ Deployment

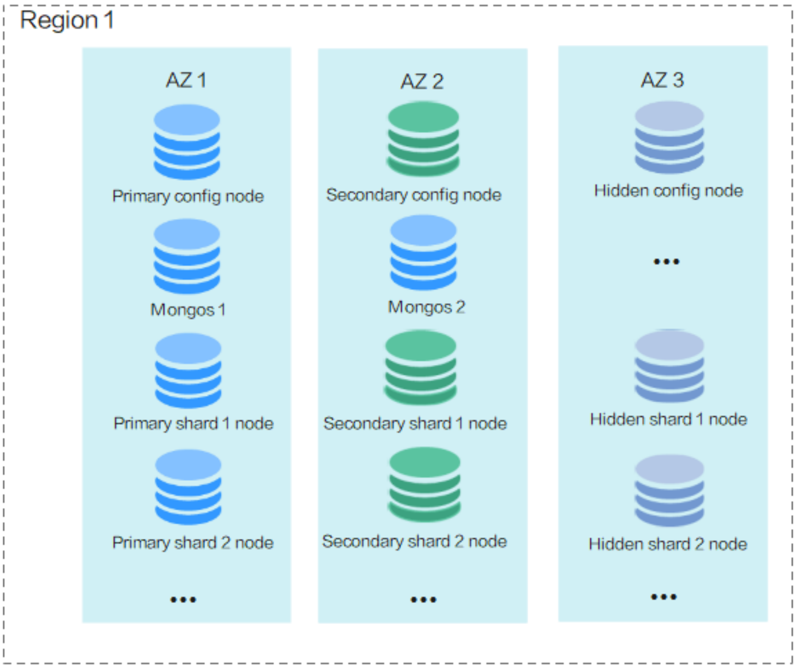

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click **Buy DB Instance**.

Step 5 Configure the instance details and click **Next**.

- **AZ**: Select three AZs as shown in **Figure 20-2**.

Figure 20-2 Selecting multiple AZs

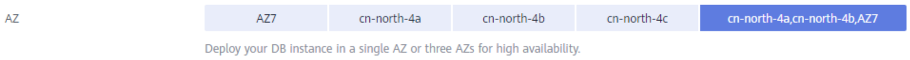

AZ: AZ7 | cn-north-4a | cn-north-4b | cn-north-4c | cn-north-4a,cn-north-4b,AZ7

Deploy your DB instance in a single AZ or three AZs for high availability.

- For details about other configuration items, see **Buying a Cluster Instance**.

Step 6 Confirm the order as prompted and complete the payment.

----End

20.2 Creating a Cross-AZ Replica Set Instance

You can deploy a replica set instance across three AZs. Multi-AZ replica set instances have higher DR capabilities than a single-AZ replica set instance and can withstand the impact caused by equipment room faults. To obtain higher DR capability, deploy resources across different AZs in the same region. If the AZ where the primary node is located fails due to power supply or network exceptions, the HA system automatically triggers a failover to ensure service continuity of a replica set instance.

This section describes how to create a replica set instance across AZs.

Precautions

- Only some regions support multi-AZ replica set instances.
- To create a multi-AZ instance, ensure that there are three or more AZs available in the region.
- If an instance is deployed in multiple AZs, the primary, secondary, and hidden nodes of the instance are deployed in three different AZs.

Deployment Architecture Comparison

- Single AZ

  If an instance is deployed in a single AZ, the primary, secondary, and hidden nodes of the instance are deployed in the same AZ.

  Figure 20-3 Single-AZ deployment

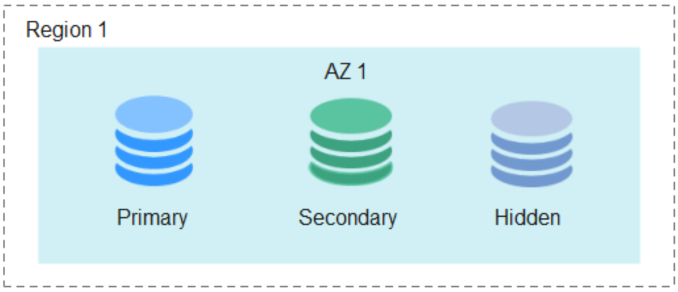

- Multiple AZs

  If an instance is deployed in multiple AZs, the primary, secondary, and hidden nodes of the instance are deployed in three different AZs for disaster recovery.

  Figure 20-4 Multi-AZ deployment

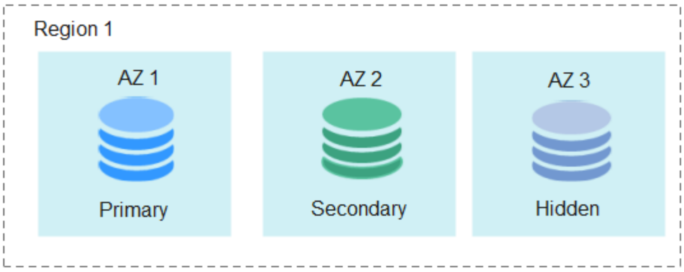

Procedure

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click **Buy DB Instance**.

Step 5 Configure the instance details and click **Next**.

- **AZ**: Select three AZs as shown in **Figure 20-5**.

Figure 20-5 Selecting multiple AZs

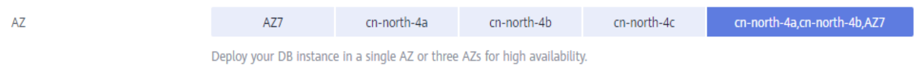

Deploy your DB instance in a single AZ or three AZs for high availability.

- For details about other configuration items, see **Buying a Replica Set Instance**.

Step 6 Confirm the order as prompted and complete the payment.

----End

21 Tags

21.1 Adding or Modifying a Tag

Tags help you identify and manage DDS resources. When there are a large number of instances, you can add tags to them to quickly filter them. An instance can be tagged during or after it is created.

This section describes how to add and modify tags after an instance is created.

Precautions

- You are advised to set predefined tags on the TMS console.
- A tag consists of a key and value. You can add only one value for each key. For details about the naming rules of tag keys and tag values, see **Table 21-1**.
- Up to 20 tags can be added for a DB instance.
- Deleting tags of a DB instance has no adverse impact on the DB instance. After all tags of a DB instance are deleted, the DB instance cannot be filtered by tag.

Procedure

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

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, click **Tags**.

Step 6 On the **Tags** page, click **Add Tag**. In the displayed dialog box, specify the tag key and value and click **OK**.

- Add a predefined tag.

  Predefined tags can be used to identify multiple cloud resources.

To tag a cloud resource, you can select a created predefined tag from the drop-down list, without entering a key and value for the tag.

For example, if a predefined tag has been created, its key is test02 and value is Project1. When you configure the key and value for a cloud resource, the created predefined tag will be automatically displayed on the page.

Figure 21-1 Adding a predefined tag

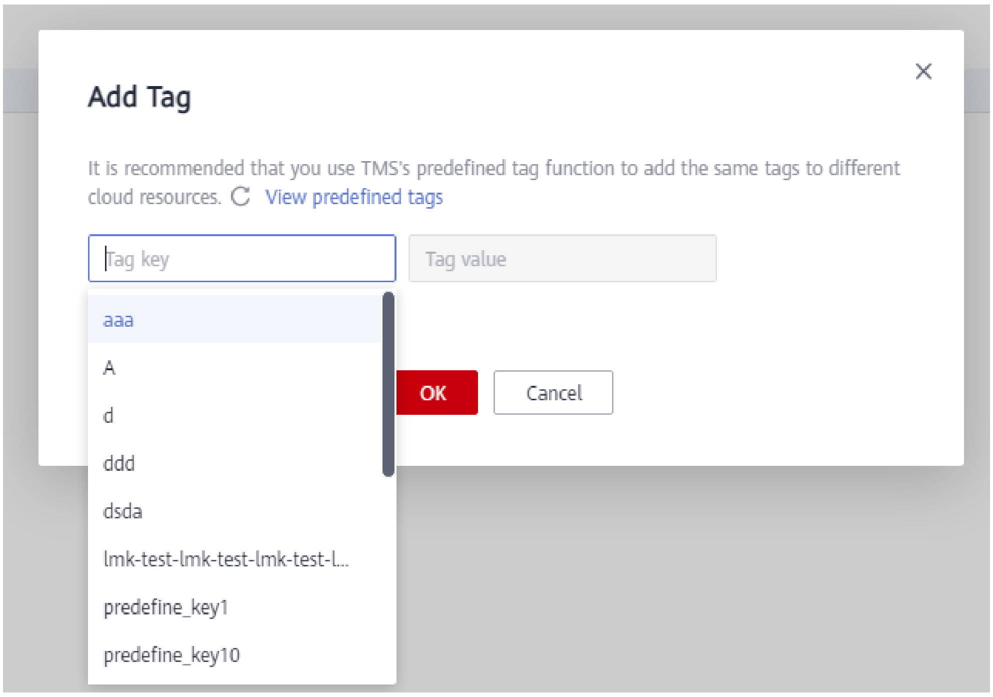

Add Tag

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

- Create a tag.

  When creating a tag, enter the tag key and value.

Figure 21-2 Adding a tag

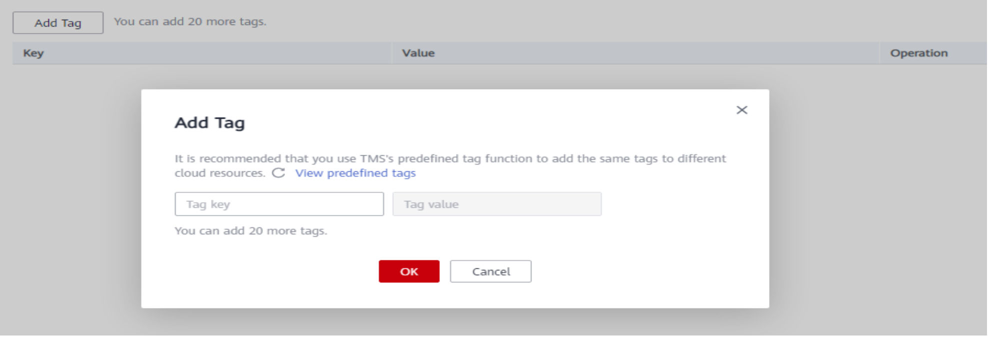

You can add 20 more tags.

Key	Value	Operation

Add Tag

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

You can add 20 more tags.

Table 21-1 Naming rules

Parameter	Requirement	Example
Tag key	– The key cannot be empty and contains 1 to 128 single-byte characters. – The key can contain UTF-8 letters (including Chinese characters), digits, spaces, and the following characters: _.:/=+-@ – Do not enter labels starting with **_sys_**, which are system labels. – The key can only consist of digits, letters, underscores (_), and hyphens (-).	Organization
Tag value	– The value can contain UTF-8 letters (including Chinese characters), digits, spaces, and the following characters: _.:/=+-@ – The value can be empty or null and contains 0 to 255 single-byte characters. – The value can only consist of digits, letters, underscores (_), periods (.), and hyphens (-).	dds_01

Step 7 View and manage tags on the **Tags** page.

You can click **Edit** in the **Operation** column to change the tag value.

NOTE

Only the tag value can be edited when editing a tag.

Figure 21-3 Tag added

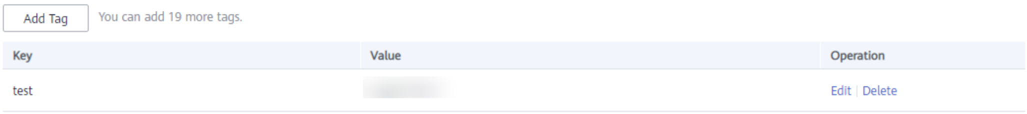

You can add 19 more tags.

Key	Value	Operation
test	[redacted]	Edit | Delete

----End

21.2 Filtering Instances by Tag

After a tag is added, you can filter instances by tag to quickly find instances of a specified category.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click **Search by Tag** in the upper right corner of the instance list.

Figure 21-4 Search by Tag

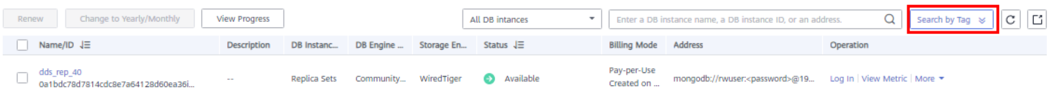

□	Name/ID	Description	DB Instanc...	DB Engine ...	Storage En...	Status	Billing Mode	Address	Operation
□	dds_rep_40 0a1bdc78d7814cdc8e7a64128d60ea36i...	--	Replica Sets	Community...	WiredTiger	Available	Pay-per-Use Created on ...	mongodb://rwuser:<password>@19...	Log In | View Metric | More

Step 5 Enter the tag key and value associated with the instance and click **Search**.

Figure 21-5 Entering the tag key and value

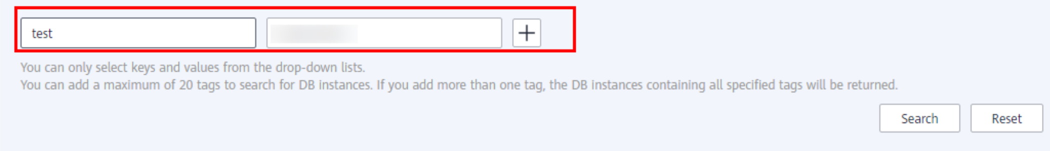
test
[redacted]

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

Step 6 View the instance information.

Figure 21-6 Viewing instance information

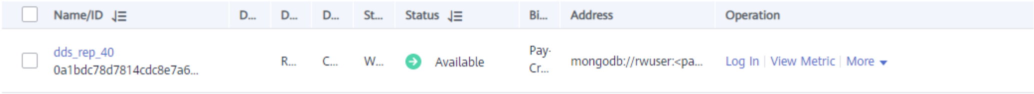

□	Name/ID	D...	D...	D...	St...	Status	Bi...	Address	Operation
□	dds_rep_40 0a1bdc78d7814cdc8e7a6...		R...	C...	W...	Available	Pay- Cr...	mongodb://rwuser:<pa...	Log In | View Metric | More

----End

21.3 Deleting a Tag

If a tag is no longer needed, you can delete the tag to unbind it from the instance.

Procedure

Step 1 Log in to the management console.

Step 2 Click in the upper left corner and select a region and a project.

Step 3 Click ☰ in the upper left corner of the page and choose **Databases** > **Document Database Service**.

Step 4 On the **Instances** page, click the instance name.

Step 5 In the navigation pane on the left, click **Tags**.

Step 6 On the **Tags** page, locate the tag to be deleted and click **Delete** in the **Operation** column. In the displayed dialog box, click **Yes**.

Figure 21-7 Deleting a tag

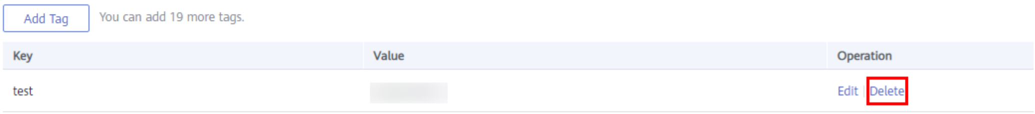

You can add 19 more tags.

Key	Value	Operation
test		Edit Delete

Step 7 After the tag is deleted, it is no longer displayed on the **Tags** page.

----End

22 Quotas

Quotas are enforced for service resources on the platform to prevent unforeseen spikes in resource usage. For example, the maximum number of DDS DB instances that can be created varies depending on the DB instance type. You can apply for increasing quotas if necessary.

This section describes how to view the usage of each type of DDS resource and the total quotas in a specified region.

Viewing Quotas

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 In the upper right corner of the DDS console, choose **Resources** > **My Quota**.

Figure 22-1 My Quota

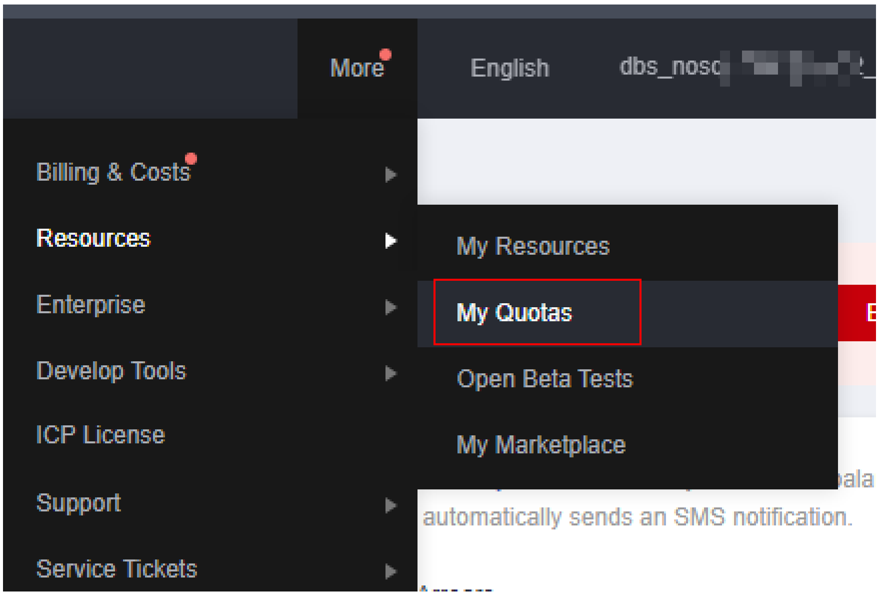

Step 4 View the used and total quota of each type of DDS resource.

----End

Increasing Quotas

Step 1 Log in to the management console.

Step 2 Click 📍 in the upper left corner and select a region and a project.

Step 3 In the upper right corner of the DDS console, choose **Resources** > **My Quota**.

Step 4 Click **Increase Quota**.

Step 5 On the **Create Service Ticket** page, configure parameters as required.

In the **Problem Description** area, fill in the content and reason for adjustment.

Step 6 After all necessary parameters are configured, select **I have read and agree to the Tenant Authorization Letter and Privacy Statement** and click **Submit**.

----End

23 DDS Usage Suggestions

23.1 Design Rules

Naming

- The name of a database object (database name, table name, field name, or index name) has to start with a lowercase letter and must be followed by a letter or digit. The length of the name cannot exceed 32 bytes.
- The database name cannot contain special characters ("".$\/*?~#:|") or null character (\0). The database name cannot be a system database name, such as **admin**, **local**, and **config**.
- The database collection name can only contain letters and underscores (_). The name cannot be prefixed with "system". The total length of <*Database name*>.<*Collection name*> cannot exceed 120 characters.

Index

You can use indexes to avoid full table scans and improve query performance.

- A column index can have up to 512 bytes, an index name can have up to 64 characters, and a composite index can have up to16 columns.
- The total length of <*Database name*>.<*Collection name*>.$<*Index name*> cannot exceed 128 characters.
- Create indexes for fields with high selectivity. If you create indexes for low selective fields, large result sets may be returned. This should be avoided.
- Write operations on a collection will trigger more I/O operations on indexes in the collection. Ensure that the number of indexes in a collection does not exceed 32.
- Do not create indexes that will not be used. Unused indexes loaded to the memory will cause a waste of memory. In addition, useless indexes generated due to changes in service logic must be deleted in a timely manner.
- Indexes must be created in the background instead of foreground.
- An index must be created for the sort key. If a composite index is created, the column sequence of the index must be the same as that of the sort key. Otherwise, the index will not be used.

- Do not create an index based on the leading-edge column of a composite index. If the leading-edge column of a composite index is the column used in another index, the smaller index can be removed. For example, a composite index based on "firstname" and "lastname" can be used for queries on "firstname". In this case, creating another firstname-based index is unnecessary.
- Creating indexes consumes a lot of I/O and compute resources. You are advised to create indexes during off-peak hours. Do not concurrently create more than five indexes. If you need to create multiple indexes for a given collection, run the **createIndexes** command to deliver multiple indexes at a time to reduce performance loss.

Sharding

You can shard collections to maximize the cluster performance. For details, see, **Sharding a Collection**.

Suggestions for sharding collections:

- In scenarios where the data volume is large (more than one million rows) and the write/read ratio is high, sharding is recommended if the data volume increases with the service volume.
- If you shard a collection using a hashed shard key, pre-splitting the chunks of the sharded collection can help reduce the impact of automatic balancing and splitting on service running.
- If sharding is enabled for a non-empty collections, the time window for enabling the balancer must be set during off-peak hours. Otherwise conflicts may occur during data balancing between shards and service performance will be affected.
- If you want to perform a sort query based on the shard key and new data is evenly distributed based on the shard key, you can use ranged sharding. In other scenarios, you can use hashed sharding.
- Properly design shard keys to prevent a large amount of data from using the same shard key, which may lead to jumbo chunks.
- If a sharded cluster is used, you must run **flushRouterConfig** after running **dropDatabase**. For details, see **How Do I Prevent dds mongos Cache Problem?**
- The update request of the service must match the shard key. When a sharded table is used, an error will be reported for the update request and "An upsert on a sharded collection must contain the shard key and have the simple collation" will be returned in the following scenarios:
  - The filter field of the update request does not contain the shard key field and the value of **multi** is **false**.
  - The set field does not contain the shard key and the value of **upsert** is **true**.

23.2 Development Rules

Database Connections

If the maximum number of mongod or dds mongos connections is reached, your client cannot connect to the DDS instances. Each connection received by mongod or dds mongos is processed by a single thread of 1 MB stack space. As the connections increase, too many threads will increase the context switching overhead and memory usage.

- If you connect to databases from clients, calculate the number of clients and the size of the connection pool configured for each client. The total number of connections cannot exceed 80% of the maximum number of connections allowed by the current instance.
- The connection between the client and the database must be stable. It is recommended that the number of new connections per second be less than 10.
- You are advised to set the connection timeout interval of the client to at least three times the maximum service execution duration.
- For a replica set instance, the IP addresses of both the primary and standby nodes must be configured on the client. For a cluster instance, at least two dds mongos IP addresses must be configured.
- DDS uses user **rwuser** by default. When you log in as user **rwuser**, the authentication database must be **admin**.

Reliability

Rules for setting write concern: For mission-critical services, set write concern to {w:n},n>0. A larger value is better consistency but poorer performance.

- **w:1** means that a confirmation message was returned after data was written to the primary node.
- **w:1,journal:true** means that the result was returned after data was written to the primary node and logs.
- **w:majority** means that the result was returned after data was written to more than half of the total standby nodes.

NOTE

If data is not written using **w:majority**, the data that is not synchronized to the standby node may be lost when a primary/standby switchover occurs.

If high reliability is required, deploy a cluster in three AZs.

Performance

Specification

- The service program is not allowed to perform full table scanning.
- During the query, select only the fields that need to be returned. In this way, the network and thread processing loads are reduced. If you need to modify

data, modify only the fields that need to be modified. Do not directly modify the entire object.

- Do not use $not. DDS does not index missing data. The $not query requires that all records be scanned in a single result collection. If $not is the only query condition, a full table scan will be performed on the collection.
- If you use $and, put the conditions with the fewest matches before other conditions. If you use $or, put the conditions with the more matches first.
- In a DB instance, the total number of databases cannot exceed 200, and the total number of collections cannot exceed 500. If the number of collections is too large, the memory may be overloaded. In addition, the performance for restarting a DB instance and performing a primary/standby switchover may deteriorate due to too many collections, which affects the high availability performance in emergencies.
- Before bringing a service online, perform a load test to measure the performance of the database in peak hours.
- Do not execute a large number of concurrent transactions at the same time or leave a transaction uncommitted for a long time.
- Before rolling out services, execute query plans to check the query performance for all query types.

Suggestions

- Each connection is processed by an independent thread in the background. Each thread is allocated with 1 MB stack memory. The number of connections should not be too large. Otherwise, too much memory is occupied.
- Use the connection pool to avoid frequent connection and disconnection. Otherwise, the CPU usage is too high.
- Reduce disk read and write operations: Reduce unnecessary upsert operations.
- Optimize data distribution: Data is sharded and hot data is distributed evenly between shards.
- Reduce lock conflicts: Do not perform operations on the same key too frequently.
- Reduce lock wait time: Do not create indexes on the frontend.

Notice

During the development process, each execution on a collection must be checked using explain() to view its execution plan. Example:

db.T_DeviceData.find({"deviceId":"ae4b5769-896f"}).explain();

db.T_DeviceData.find({"deviceId":"77557c2-31b4"}).explain("executionStats");

A covered query does not have to read a document and returns a result from an index, so using a covered query can greatly improve query efficiency. If the output of explain() shows that indexOnly is true, the query is covered by an index.

Execution plan parsing:

1. Check the execution time. The smaller the values of the following parameters, the better the performance: **executionStats.executionStages.executionTimeMillisEstimate** and **executionStats.executionStages.inputStage. executionTimeMillisEstimate**

- **executionStats.executionTimeMillis** specifies how much time the database took to both select and execute the winning plan.
- **executionStats.executionStages.executionTimeMillisEstimate** specifies the execution completion time of the execution plan.
- **executionStats.executionStages.inputStage. executionTimeMillisEstimate** specifies the execution completion time of the sub-phase of the execution plan.

2. Check the number of scanned records. If the three items are the same, the index is best used.
   - **executionStats. nReturned** is the number of documents that match the query condition.
   - **executionStats .totalKeysExamined** indicates the number of scanned index entries.
   - **executionStats .totalDocsExamined** indicates the number of scanned document entries.
3. Check the stage status. The following combinations of stages can provide good performance.
   - Fetch+IDHACK
   - Fetch+ixscan
   - Limit+ (Fetch+ixscan)
   - PROJECTION+ixscan

Table 23-1 Status description

Status Name	Description
COLLSCAN	Full table scan
SORT	In-memory sorting
IDHACK	_id-based query
TEXT	Full-text index
COUNTSCAN	Number of unused indexes
FETCH	Index scanning
LIMIT	Using Limit to limit the number of returned records
SUBPLA	$or query stage without using an index
PROJECTION	Restricting the return of stage when a field is returned.
COUNT_SCAN	Number of used indexes

Cursor Usage Rules

If a cursor is inactive for 10 minutes, it will be automatically closed. You can also manually close it to save resources.

Rules for Using Distributed Transactions in Version 4.2

- Spring Data MongoDB does not support the retry mechanism after a transaction error is reported. If the client uses Spring Data MongoDB as the client to connect to MongoDB, you need to use Spring Retry to retry the transaction based on the references of Spring Data MongoDB.
- The size of the distributed transaction operation data cannot exceed 16 MB.

Precautions for Backups

Do not perform DDL operations during the backup to avoid backup failures.