# **Distributed Cache Service**

# **User Guide**

 Issue
 01

 Date
 2025-04-27





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

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

# **Trademarks and Permissions**

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

# Notice

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

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

# Huawei Cloud Computing Technologies Co., Ltd.

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

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

# **Contents**

1 Process of Using DCS	1
2 Creating a User and Granting DCS Permissions	3
3 Buying a DCS Redis Instance	7
4 Accessing a DCS Redis Instance	
4.1 Configuring Redis Network Connections	16
4.1.1 Network Conditions for Accessing DCS Redis	
4.1.2 Enabling Public Access to Redis and Obtaining the Access Addresses	17
4.2 Controlling DCS Redis Access	
4.2.1 Configuring DCS Redis Access Whitelist	20
4.2.2 Configuring a Redis Password	22
4.2.3 Transmitting DCS Redis Data with Encryption Using SSL	
4.2.4 Configuring DCS Redis ACL Users	25
4.3 Connecting to Redis on a Client	
4.3.1 Connecting to Redis on redis-cli	27
4.3.2 Connecting to Redis on Jedis (Java)	
4.3.3 Connecting to Redis on Lettuce (Java)	
4.3.4 Connecting to Redis on Redisson (Java)	
4.3.5 Connecting to Redis on redis-py (Python)	66
4.3.6 Connecting to Redis on go-redis (Go)	
4.3.7 Connecting to Redis on hiredis (C++)	71
4.3.8 Connecting to Redis on StackExchange.Redis (C#)	73
4.3.9 Connecting to Redis on phpredis (PHP)	75
4.3.10 Connecting to Redis on predis (PHP)	77
4.3.11 Connecting to Redis on ioredis (Node.js)	79
4.4 Connecting to Redis on the Console	
4.5 Public Access to a DCS Redis 3.0 Instance (Discontinued)	83
4.5.1 Enabling Public Access of a DCS Redis 3.0 Instance	
4.5.2 Connecting to Redis 3.0 over a Public Network on redis-cli	
5 Accessing a DCS Memcached Instance (Discontinued)	94
5.1 Configuring a Memcached Password	94
5.2 Connecting to Memcached on a Client	95
5.2.1 Connecting to Memcached on the Telnet	95

5.2.2 Connecting to Memcached on the Spymemcached (Java)	
5.2.3 Connecting to Memcached on the Python-binary-memcached (Python)	
5.2.4 Connecting to Memcached on the Libmemcached (C++)	
5.2.5 Connecting to Memcached on the Libmemcached (PHP)	
6 Managing Instances	110
6.1 Viewing and Modifying Basic Settings of a DCS Instance	110
6.2 Viewing DCS Background Tasks	115
6.3 Viewing Client Connection Information of a DCS Instance	
6.4 Modifying Configuration Parameters of a DCS Instance	118
6.5 Configuring DCS Instance Parameter Templates	144
6.5.1 Viewing a Parameter Template of a DCS Instance	144
6.5.2 Creating a Custom Parameter Template for a DCS Instance	169
6.6 Configuring DCS Instance Tags	195
6.7 Renaming Critical Commands for DCS Instances	197
6.8 Returning the Real IP Addresses of a Client to DCS (IP Pass-through)	198
6.9 Exporting a DCS Instance List	199
6.10 Performing a Master/Standby Switchover for a DCS Instance	199
6.11 Managing DCS Instance Shards and Replicas	
7 Backing Up or Restoring Instance Data	
7.1 DCS Backup and Restoration Overview	
7.2 Backing up DCS Instances Automatically	205
7.3 Backing up DCS Instances Manually	207
7.4 Restoring DCS Instances	208
7.5 Downloading DCS Instance Backup Files	
8 Changing an Instance	213
8.1 Modifying DCS Instance Specifications	213
8.2 Adjusting DCS Instance Bandwidth	
8.3 Changing Cluster DCS Instances to be Across AZs	
8.4 Upgrading Minor or Proxy Versions of a DCS Instance	229
9 Managing Lifecycle of an Instance	232
9.1 Restarting a DCS Instance	
9.2 Starting or Stopping a DCS Instance	
9.3 Deleting a DCS Instance	233
9.4 Clearing DCS Instance Data	235
10 Diagnosing and Analyzing an Instance	236
10.1 Querying Big Keys and Hot Keys in a DCS Redis Instance	
10.2 Scanning and Deleting Expired Keys in a DCS Redis Instance	
10.3 Analyzing Redis Backup Offline	
10.4 Diagnosing a DCS Redis Instance	247
10.5 Viewing Slow Queries of a DCS Redis Instance	248

10.6 Viewing Redis Run Logs	
10.7 Viewing Audit Logs of a DCS Redis Instance	250
11 Migrating Instance Data	
11.1 DCS Data Migration Overview	253
11.2 Migration Solution Notes	
11.3 Migrating Data Between DCS Instances	
11.3.1 Online Migration Between Instances	
11.3.2 Backup Import Between DCS Redis Instances	
11.4 Migrating Data from Self-Hosted Redis to DCS	273
11.4.1 Migrating Self-Built Redis Online	273
11.4.2 Self-Hosted Redis Migration with Backup Files	
11.4.3 Self-Hosted Redis Migration with redis-cli (AOF)	
11.4.4 Self-Hosted Redis Migration with redis-cli (RDB)	
11.4.5 Self-Hosted Redis Cluster Migration with redis-shake (Online)	285
11.4.6 Self-Hosted Redis Cluster Migration with redis-shake (RDB)	
11.5 Migration from Another Cloud	290
11.5.1 Migrating Redis from Another Cloud Online	
11.5.2 Backup Import from Another Cloud	294
11.5.3 Online Migration from Another Cloud Using Rump	297
11.5.4 Migrating from Another Cloud Online Using redis-shake	299
11.5.5 Backup Import from Another Cloud Using redis-shake	305
12 Testing Instance Performance	307
12.1 Testing Redis Performance Using memtier_benchmark	
12.2 Testing Redis Performance Using redis-benchmark	
12.3 Comparing redis-benchmark and memtier_benchmark	313
12.4 Reference for a Redis Performance Test	313
12.4.1 Test Data of Master/Standby DCS Redis 3.0 Instances	313
12.4.2 Test Data of Proxy Cluster DCS Redis 3.0 Instances	
12.4.3 Test Data of Master/Standby DCS Redis 4.0 or 5.0 Instances	
12.4.4 Test Data of Proxy Cluster DCS Redis 4.0 or 5.0 Instances	317
12.4.5 Test Data of Redis Cluster DCS Redis 4.0 or 5.0 Instances	
12.4.6 Test Data of Master/Standby DCS Redis 6.0 Instances	321
12.4.7 Test Data of Redis Cluster DCS Redis 6.0 Instances	323
12.4.8 Test Data of Redis Backup, Restoration, and Migration	
13 Applying for More DCS Quotas	330
14 Viewing Monitoring Metrics and Configuring Alarms	332
14.1 DCS Metrics	
14.2 Common DCS Metrics	
14.2 Common DCS Metrics 14.3 Viewing DCS Metrics	

15 Viewing DCS Audit Lo	ogs391
-------------------------	--------

# Process of Using DCS

# How to Manage DCS Instances

You can access Distributed Cache Service (DCS) from the web-based management console or by using RESTful application programming interfaces (APIs) through HTTPS requests.

• Using the console

You can sign up and log in to the **console**, click in the upper left corner on the homepage, and choose **Middleware** > **Distributed Cache Service for Redis**.

For details on how to use the DCS console, see chapters from **Buying a DCS Redis Instance** to **Migrating Instance Data**.

DCS monitoring data is recorded by Cloud Eye. To view the monitoring metrics or configure alarm rules, go to the Cloud Eye console. For details, see **Viewing DCS Metrics**.

If you have enabled Cloud Trace Service (CTS), DCS instance operations are recorded by CTS. You can view the operations history on the CTS console. For details, see **Viewing DCS Audit Logs**.

• Using APIs

DCS provides RESTful APIs for you to integrate DCS into your own application system. For details about DCS APIs and API calling, see the **Distributed Cache Service API Reference**.

### NOTICE

- 1. For DCS instance functions with open APIs, manage them using the console or calling the APIs. For those without open APIs, manage them using the console.
- 2. For details about APIs for monitoring and auditing, see the **Cloud Eye** and **Cloud Trace Service (CTS)** documentation.

# Using DCS



# Figure 1-1 Process of using DCS

- 1. Creating a User and Granting DCS Permissions
- 2. Buying a DCS Redis Instance
- 3. Accessing a DCS Redis Instance

Redis instances can be accessed on a client or the DCS console.

4. Managing DCS Instances and Data

Learn about Managing Instances, Managing Lifecycle of an Instance, Changing an Instance, Testing Instance Performance, Diagnosing and Analyzing an Instance, Backing Up or Restoring Instance Data, and Migrating Instance Data.

# **2** Creating a User and Granting DCS Permissions

This section describes how to use **Identity and Access Management (IAM)** to implement fine-grained permissions control for your DCS resources. With IAM, you can:

- Create IAM users for employees based on your enterprise's organizational structure. Each IAM user will have their own security credentials for accessing DCS resources.
- Manage permissions on a principle of least permissions (PoLP) basis.
- Entrust a Huawei Cloud account or cloud service to perform efficient O&M on your DCS resources.

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

This section describes the procedure for granting the **DCS ReadOnlyAccess** permission (see **Figure 2-1**) as an example.

# Prerequisites

Learn about the permissions (see **System-defined roles and policies supported by DCS**) supported by DCS and choose policies or roles according to your requirements. For the permissions of other services, see **Permissions Policies**.

# **Process Flow**



# Figure 2-1 Process of granting DCS permissions

### 1. Create a user group and assign permissions.

Create a user group on the IAM console, and assign the **DCS ReadOnlyAccess** policy to the group.

2. Create an IAM user.

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 DCS console using the newly created user, and verify that the user only has read permissions for DCS.

- Choose Distributed Cache Service in Service List. Then click Buy DCS Instance in the upper right corner of the DCS console. If a DCS instance cannot be purchased, the DCS ReadOnlyAccess policy has already taken effect.
- Choose any other service in Service List. If a message appears indicating that you have insufficient permissions to access the service, the DCS ReadOnlyAccess policy has already taken effect.

# **DCS Custom Policies**

You can create custom policies to supplement the system-defined policies of DCS. For the actions that can be added to custom policies, see **Permissions Policies and Supported Actions**.

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

• Visual editor: Select cloud services, actions, resources, and request conditions. This does not require knowledge of policy syntax.

• JSON: Create a JSON policy or edit an existing one.

For details, see **Creating a Custom Policy**. The following section contains examples of common DCS custom policies.

#### **NOTE**

{

}

Due to data caching, a policy involving OBS actions will take effect five minutes after it is attached to a user, user group, or project.

# **Example Custom Policies**

• Example 1: Grant permission to delete and restart DCS instances and clear data of an instance.

• Example 2: Grant permission to deny DCS deletion.

A policy with only "Deny" permissions must be used together with other policies. If the permissions granted to an IAM user contain both "Allow" and "Deny", the "Deny" permissions take precedence over the "Allow" permissions.

Assume that you want to grant the permissions of the **DCS FullAccess** policy to a user but want to prevent them from deleting DCS instances. You can create a custom policy for denying DCS deletion, and attach this policy together with the **DCS FullAccess** policy to the user. As an explicit deny in any policy overrides any allows, the user can perform all operations on DCS instances excepting deleting them. Example policy denying DCS deletion:

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

# **DCS Resources**

A resource is an object that exists within a service. DCS resources include instance. To select these resources, specify their paths.

Resource	Resource Name	Path
instance	Instance	[Format] DCS:*:*: instance: <i>instance ID</i>
		[Note]
		For instance resources, DCS automatically generates the prefix ( <b>DCS:*:*:instance:</b> ) of the resource path.
		For the path of a specific instance, add the instance ID to the end. You can also use an asterisk * to indicate any instance. For example:
		DCS:*:*:instance:* indicates any DCS instance.

# Table 2-1 DCS resources and their paths

# **3** Buying a DCS Redis Instance

You can buy one or more DCS Redis instances with the required computing capabilities and storage space based on service requirements.

# **Preparing Required Resources**

DCS Redis instances are deployed in Virtual Private Clouds (VPCs), and bound to specific subnets. In this way, Redis instances are isolated with virtual networks that users can manage by themselves. To use professional edition Redis, security groups are required.

Therefore, before creating a Redis instance, prepare a VPC, security group, and subnet if you do not have one yet.

Resource	Requirement	Operations
VPC and subnet	<ul> <li>Different Redis instances can use the same or different VPCs and subnets as required. Note the following when creating a VPC and subnet:</li> <li>The VPC and the DCS instance must be in the same region.</li> <li>Retain the default settings uplace otherwise specified</li> </ul>	For details on how to create a VPC and subnet, see <b>Creating</b> <b>a VPC</b> . If you need to create and use a new subnet in an existing VPC, see <b>Creating a</b> <b>Subnet for the VPC</b> .

<b>Tuble 5</b> T Dependency resources of a Des instance
---

Resource	Requirement	Operations
Security group (Only profession al edition DCS Redis instances require security groups.)	<ul> <li>Different Redis instances can use the same security group or different security groups. Note the following when creating a security group:</li> <li>Set <b>Template</b> to <b>Custom</b>.</li> <li>After a security group is created, retain the default inbound and outbound rules.</li> <li>To use professional edition DCS Redis 6.0, you must add the security group rules described in <b>Table 3-2</b>. You can also add other rules as</li> </ul>	For details on how to create a security group, see <b>Creating a</b> <b>Security Group</b> . For details on how to add rules to a security group, see <b>Adding a Security</b> <b>Group Rule</b> .

Table 3-2 Security group rules

Direction	Protocol	Port	Source
Inbound	ТСР	6379	0.0.0/0

# **Buying a DCS Redis Instance**

- **Step 1** Go to the **Buy DCS Instance** page.
- Step 2 Choose a purchase mode. Quick Config and Custom are available.

In **Quick Config**, several common **Specification Settings** are available, as shown in **Figure 3-1**. If none of them meet your needs, go to the **Custom** tab page to customize instance types and specifications, as shown in **Figure 3-2**.

**NOTE** 

Quick Config and Custom differ only in specification selection.

# Figure 3-1 Quick Config

sic Settings							
ng Mode	Yearly/Monthly	Pay-per-use					
ion	•		~				
	Regions are geographic a	reas isolated from each other. R	Resources are region-specifi	ic and cannot be used across re	gions through internal netwo	rk connections. For low netwo	rk latency and quick resource ac
ect			~				
ecification Settings		_					
Basic - 16	6GB and some data	Basic - 320	GB	Basic - 32	GB	For high performance an accelerations	onal - 8GB
Basic - 16	6GB and some data Redis 6.0	Basic - 320 For general service accele	GB rration Redis 6.0	For general service accele	GB aration Redis 6.0	For high performance an accelerations	onal - 8GB ad concurrency in all Redis 6.0
Basic - 16 For development testing acceleration Version	6GB and some data Redis 6.0 Master/Standby	For general service accele Version	GB rration Redis 6.0 Proxy Cluster	For general service accele Version Instance Type	GB rration Redis 6.0 Redis Cluster	For high performance an accelerations Version	onal - 8GB id concurrency in all Redis 6.0 Master/Standby
Basic - 16 Basic - 16 For development testing acceleration Version	BGB and some data Redis 6.0 Master/Standby x86	For general service accele Version Instance Type CPU Architecture	GB realise 6.0 Proxy Cluster x86	For general service accele Version Instance Type CPU Architecture	CB Redis 6.0 Redis Cluster x86	Profession For high performance an accelerations Version Instance Type CPU Architecture	nal - 8GB Id concurrency in al Redis 6.0 Master/Standby x86
Basic - 16 For development testing Acceleration Instance Type CPU Architecture Replicas	SGB and some data Redis 6.0 Master/Standby x86 2	For general service accele Version Instance Type CPU Architecture Replicas	CB reation Proxy Cluster x86 2	For general service accele Version Instance Type CPU Architecture Replicas	SB reation Redis 6.0 Redis Cluster x86 2	Profession For high performance an accelerations Version Instance Type CPU Architecture Replicas	onal - 8GB d concurrency in al Redis 6.0 Master/Standby x86 2

Figure 3-2 Custom

Quick Config Custon	n
Basic Settings	
Billing Mode	Yearly/Monthly Pay-per-use
Region	9 v
	Regions are geographic areas isolated from each other. Resources are region-specific and cannot be used across regions through internal network connections. For low network latency and quick resource access, select the nearest region.
Project	v
0	
Specification Setting:	S
Cache Engine	Reds
Edition (?)	Basit Professional
	Compatible with Redis 6.0, with powerful litered storage (SSD) of cold/hot data. Meets the demands of online or heavy services.
CPU Architecture	196 Am
010110100010	
Version ③	60 50 40
	Fixed on instance creation. To change to a higher version, buy another instance and migrate data to it.
Instance Type  )	Single-node Master/Standay Read/Write splitting Proxy Cluster Redis Cluster
	Bachup   Falover   Persistence

- Step 3 Specify Billing Mode.
- **Step 4** Select a region closest to your application to reduce latency and accelerate access.
- **Step 5** Select a project. By default, each region corresponds to a project.
- **Step 6** Configure specifications. If you choose **Quick Config**, select one of the common specifications by referring to **Table 3-3**. If you choose **Custom**, configure instance specifications by referring to **Table 3-4**.

ltem	Description
Basic - memory/ Professional -	Basic and Professional editions are available. Comparing Professional and Basic Editions
memory	For example, <b>Basic - 16GB</b> is a basic edition instance with 16 GB memory.
Version	Version of the Redis instance. For details, see <b>Comparing</b> <b>Redis Versions</b> .
	The Redis version cannot be changed once the instance is created. To use a later Redis version, create another DCS Redis instance and then migrate data from the old instance to the new one.
Instance Type	Master/Standby, Redis Cluster, and Proxy Cluster instances can be quickly configured. For details, see <b>DCS Instance Types</b> .
CPU Architecture	x86-based CPU can be quickly configured.
Replicas	Two replicas mean that the instance has two nodes (one master and one standby).
AZ	AZs where the master node and standby node of the instance are located.

Table 3-3 Specifications (Quick Config)

ltem	Description		
Cache Engine	Only Redis is available.		
Edition	Basic and Professional editions are available. Comparing Professional and Basic Editions		
	<b>Basic</b> and <b>Professional</b> are available only when <b>6.0</b> is selected. Other versions support only <b>Basic</b> .		
Sub-edition	<b>Performance</b> and <b>Storage</b> are available for the professional edition.		
	Sub-edition is required only when <b>Professional</b> is selected.		
CPU Architecture	<ul> <li>x86 and Arm are available.</li> <li>x86 is recommended. Arm is unavailable in some regions</li> </ul>		

Item	Description		
Version	Currently supported Redis versions: 4.0, 5.0, 6.0, and 7.0. Currently, DCS for Redis 7.0 is in open beta test (OBT). If it is unavailable on the console, contact customer service to add you to the whitelist. For details about the differences between Redis versions, see <b>Comparing Redis Versions</b> .		
	The Redis version cannot be changed once the instance is created. To use a later Redis version, create another DCS Redis instance and then migrate data from the old instance to the new one.		
Instance Type	Single-node, master/standby, Proxy Cluster, Redis Cluster, and read/write splitting types are supported. For more information, see <b>DCS Instance Types</b> .		
	The master/standby type is available in the Redis 6.0 enterprise edition.		
	The supported Redis versions and instance types vary across regions.		
AZ	AZ: If Instance Type is master/standby, read/write splitting, Proxy Cluster, or Redis Cluster, AZ and Standby AZ are displayed. In this case, select AZs for the master and standby nodes of the instance.		
	Each region consists of multiple AZs with physically isolated power supplies and networks. The master and standby nodes of a master/standby, read/write splitting, o cluster DCS instance can be deployed in different AZs. Applications can also be deployed across AZs to achieve high availability (HA) for both data and applications		
	NOTE		
	• If instance nodes in an AZ are faulty, nodes in other AZs will not be affected. This is because when the master node is faulty, the standby cache node will automatically become the master node to provide services. Such deployment achieves better disaster recovery.		
	• Deploying a DCS instance across AZs slightly reduces network efficiency compared with deploying an instance within an AZ. Therefore, if a DCS instance is deployed across AZs, synchronization between master and standby cache nodes is slightly less efficient.		
	<ul> <li>To accelerate access, deploy your instance and your application in the same AZ.</li> </ul>		
	• To purchase a DCS instance in the CN South-Guangzhou region, AZ 1–5 and 6–7 cannot be selected as the master and standby AZs for an instance. The reason is that AZ 1–5 are physically distant from AZ 6–7 which may cause high internal latency. That is, if one of AZs 1–5 is selected as the primary AZ, AZs 6–7 cannot be selected as the standby AZ. This restriction applies only in the CN South-Guangzhou region.		

ltem	Description		
Replicas	Enter the number of replicas. Replicas are DCS instance nodes. One replica indicates only a master node. Two replicas indicate a master node and a standby node. Three replicas indicate a master node and two standby nodes.		
	Replica quantity range varies by version and instance type. <b>Replicas</b> cannot be set for single-node instances.		
Sharding	This parameter is available only for cluster instances, ar is single-choice. The shard size and quantity cannot be specified at once.		
	• <b>Use default</b> : You do not need to specify the shard size and quantity. Select one of the default instance specifications.		
	• <b>Custom shard size</b> : Select a shard size. Then select an instance specification.		
	• <b>Custom shard quantity</b> : Select a shard quantity. Then select an instance specification.		
Instance Specification	In the <b>Instance Specification</b> area, select memory as required. For more information about the instance performance, see <b>DCS Instance Specifications</b> . The default memory quota is displayed on the console.		
	To increase quota, click <b>Increase quota</b> below the specifications. On the displayed page, fill in a quota application form and click <b>Submit</b> .		

**Step 7** Configure instance network settings.

- 1. Select the created VPC and Subnet.
  - To access the instance in an Elastic Cloud Server (ECS), select the VPC where the ECS is.
  - The VPC and subnet are fixed once the DCS instance is created.
  - A shared VPC implements network resource sharing, and unified and efficient management and control at low O&M costs.
- 2. In the IPv4 Address area, set the instance (private) IP address.

Redis Cluster and professional edition instances only support automaticallyassigned IP addresses. The other instance types support both automaticallyassigned IP addresses and manually-specified IP addresses. You can manually specify an IP address for your instance as required.

3. Configure **Port**. For basic edition Redis instances, you can specify a port numbering in the range from 1 to 65535. If no port is specified, the default port 6379 will be used.

For Redis 6.0 professional, you cannot customize a port. The default port 6379 will be used.

4. Select the created security group from the **Security Group** drop-down list.

A security group is a set of rules that control access to ECSs. It provides access policies for mutually trusted ECSs with the same security protection requirements in the same VPC.

Only professional edition DCS Redis instances require security groups. Basic edition DCS Redis instances are based on VPC endpoints and do not support security groups. To control access to these instances, **configure a whitelist** after the instances are created.

If port 6379 is not enabled for the selected security group, the **Enable port 6379** check box is displayed and selected by default, indicating that after the instance is created, port 6379 will be enabled for the selected security group. If port 6379 is not enabled for the selected security group, connections to the instance may fail.

## Step 8 Set Instance Name.

When you create only one instance at a time, the value of **Name** can contain 4 to 64 characters. When you create more than one instance at a time, the value of **Name** can contain 4 to 56 characters. These instances are named in the format of "*name-n*", in which *n* starts from 000 and is incremented by 1. For example, if you create two instances and set **Name** to **dcs\_demo**, the two instances are respectively named as **dcs\_demo-000** and **dcs\_demo-001**.

**Step 9** Specify **Enterprise Project**. An enterprise project manages cloud resources by gathering relevant ones together.

If you cannot select an enterprise project, check your permissions. For details, see Why Can't I Select the Required Enterprise Project When Creating a DCS Instance?

- **Step 10** Set the instance password.
  - Select Yes or No for Password Protected.

**NOTE** 

- Password-free access carries security risks. Exercise caution when selecting this mode.
- After creating a password-free DCS Redis instance, you can set a password for it later by using the password reset function. For details, see Changing Password Settings for DCS Redis Instances.
- Password and Confirm Password: These parameters indicate the password of accessing the DCS Redis instance, and are displayed only when Password Protected is set to Yes.

### **NOTE**

- For security purposes, if password-free access is disabled, the system prompts you to enter an instance-specific password when you are accessing the DCS Redis instance.
- Keep your instance password secure and change it periodically. The system cannot detect your password.
- Step 11 Click Advanced Settings and set the following information as required.
  - 1. Configure **Parameter Configuration**. Retain **Default templates** or select **Use custom template** as required.

If you select **Use custom template**, select one from the drop-down list box. To view or modify the configuration in the selected template, click **View**  **Parameter**. If no custom parameter template of the selected instance version and type is available, the selection box is empty. In this case, click **View Parameter Templates** to go to the template creation page to create a template. For details, see **Creating a Custom Parameter Template for a DCS Instance**.

Parameter Configuration	Default templates	Custom template	
		✓ Q	View Parameter Templates

2. To configure instance backup policies, enable **Auto Backup**.

This parameter is displayed only when the instance type is master/standby, read/write splitting, or cluster. For details about instance backup and backup policies, see **Backing Up and Restoring Instances**.

- 3. Currently, **Public Access** cannot be configured during instance creation. Configure it later on the instance details page. For details, see **Enabling Public Access to Redis and Obtaining the Access Addresses**.
- 4. Add a tag.

Tags are used to identify cloud resources. When you have many cloud resources of the same type, you can use tags to classify cloud resources by dimension (for example, by usage, owner, or environment).

If tag policies for DCS have been set in your organization, add tags to DCS instances based on these policies. If a tag does not comply with the tag policies, DCS instance creation may fail. Contact your organization administrator to learn more about tag policies.

- If you have created predefined tags, select a predefined pair of tag key and value. Click View predefined tags. On the Tag Management Service (TMS) console, view predefined tags or create new tags.
- You can also add a tag by entering the tag key and value. For details about how to name tags, see Managing Tags.
- 5. Rename critical commands.

Currently, you can only rename the **COMMAND**, **KEYS**, **FLUSHDB**, **FLUSHALL**, **HGETALL**, **SCAN**, **HSCAN**, **SSCAN**, and **ZSCAN** commands. For Proxy Cluster instances, you can also rename the **DBSIZE** and **DBSTATS** commands.

6. Specify the maintenance window.

Choose a window for DCS O&M personnel to perform maintenance on your instance. You will be contacted before any maintenance activities are performed.

- 7. Enter a description of the instance.
- **Step 12** Specify **Required Duration**. Determine the purchase duration and whether to enable auto-renewal only when purchasing a yearly/monthly Redis instance.
- Step 13 Specify Quantity.
- Step 14 Click Next.

The displayed page shows the instance information you have specified.

- **Step 15** Confirm the instance information and submit the request.
- **Step 16** After the task is successfully submitted, the **Cache Manager** page is displayed. When the new instance is in the **Running** state, the instance is created successfully.

----End

# **4** Accessing a DCS Redis Instance

# 4.1 Configuring Redis Network Connections

# 4.1.1 Network Conditions for Accessing DCS Redis

You can access a DCS instance through any Redis client. For details about Redis clients, see the **Redis official website**.

There are different constraints when a client connects to Redis in certain ways:

• Accessing a Redis instance on a client within the same VPC

The ECS where the client is installed must be in the same VPC as the DCS Redis instance. An ECS and a DCS instance can communicate with each other only when they belong to the same VPC. Redis 3.0 and 6.0 enterprise: The instance and the ECS must either be configured with the same security group or use different security groups but can communicate with each other as configured by the security group rules. Redis 4.0 and later: The IP address of the ECS must be on the whitelist of the DCS instance.

For details about how to configure a security group, see **How Do I Configure a Security Group?** For details about how to configure a whitelist, see **Managing IP Address Whitelist**.

• Accessing a Redis instance on a client across VPCs in the same region

If the client and DCS Redis instance are not in the same VPC, connect them by establishing a VPC peering connection. For details, see **Does DCS Support Cross-VPC Access?** 

• Accessing a Redis instance of another region on a client

If the client server and the Redis instance are not in the same region, connect the network using Direct Connect. For details, see **What Is Direct Connect**.

To access a Redis instance across regions, the instance domain names cannot be resolved across regions. Therefore, the instance cannot be accessed at its domain name addresses. You can manually map the domain name to the IP address in the **hosts** file, or access the instance at its IP address.

• Public network access

To access a Redis 4.0 or later instance on a client over a public network, enable public access for the instance by referring to **Enabling Public Access to Redis and Obtaining the Access Addresses**.

To access a Redis 3.0 instance over a public network on a client, the instance needs to be configured with security rules. If SSL encryption is disabled, allow public access through port 6379. If SSL encryption is enabled, allow public access through port 36379. For details, see **the "How Do I Configure a Security Group?" FAQ**.

# 4.1.2 Enabling Public Access to Redis and Obtaining the Access Addresses

Public access to Redis 4.0 and later instances can be enabled using Elastic Load Balance (ELB). This section describes how to enable the public access, obtain the access addresses and ports, and add the instances or load balancers to IP whitelists. To enable public access to Redis 3.0 instances, see **Public Access to a DCS Redis 3.0 Instance (Discontinued)**.

# Notes and Constraints

- Public access to Redis instances may be unavailable in some cases. Contact customer service to enable public access whitelists.
- Public access can be enabled for single-node, master/standby, read/write splitting, and Proxy Cluster instances.
- Public access to Redis on a client has higher network latency than access to Redis on a client within a VPC.
- SLA does not include client access exceptions caused by public network performance issues.

# Prerequisites

• A load balancer has been prepared. Learn how to create one by referring to **Creating a Dedicated Load Balancer**. The load balancer must meet the following requirements.

**NOTE** 

- The type is **Dedicated** and **IP as a Backend** is enabled.
- Network load balancing(TCP/UDP) is enabled in the specifications.
- The VPC is the same as the Redis instance.
- An Elastic IP (EIP) is bound.
- A port is available.
- Binding a load balancer to multiple DCS instances will limit the Redis performance to balancer specifications.
- For network security purposes, the Redis instance must be passwordprotected. Password-free instances do not support public access. To change the instance to be password-protected, see **Resetting an Instance Password**.

# Procedure

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

- **Step 2** Click <sup>Q</sup> in the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click an instance name to go to the instance basic information page.
- Step 5 Click On after Public Access.
- **Step 6** Select desired balancers and click **OK**.

If there is no available load balancer, click **Create one** to go to the ELB console. If the load balancer exists, but is not in the list, check whether the load balancer can be bound by referring to **Prerequisites**.

#### **NOTE**

- When a load balancer is bound to a Redis instance, do not delete the load balancer and listener. Ensure that the load balancer is available or public access to Redis may be affected.
- To delete a load balancer, unbind it (disable public access) on the Redis instance details page. Then, delete it on the ELB console.

Figure 4-1 Binding a load balancer

Balancer	Name/ID \ominus		IP Addresses ⇔		
	elt a3	)e78-b4dd6c313767	1: (private IPv4) 1: iblic IPv4)		
	elt 3e	>5e0-74ff996e16be □	1: ivate IPv4) 1: iblic IPv4)		
	O elt 5d	of06-37ef85713c08	1:     vrivate IPv4)       1:     vic IPv4)		
	elt ed	la3-e23cdb6b26f5	1: ivate IPv4)		

- **Step 7** When the task for enabling public access is in the **Successful** state, the public access is enabled.
- **Step 8** Choose **Basic Information** in the navigation pane and check public access. To disable public access, click **Off**.
  - The **EIP** is the public access address of the Redis instance and the port of the **Listener** is the public access port.

## Figure 4-2 Public access address

Connection ⑦		
Password Protected	Yes	
Connection Address	redis-19c10e	3.com:6379 🗇 🖉
IP Address	192.16 7:6379 🗇 🖉	
Public Access	On Off	
	Load balancer:elb- EIP: 10 86 7 Listener: listener-19c10e	0a1b 🕐 Port <mark> 6379</mark>

• Enabling public access for a master/standby instance generates two listeners. The listener (starting with listener-master) listens to the master node. The listener (starting with listener-slave) listens to the standby node. For public access to a master/standby instance, use the master listener port to connect to the master node of the instance. To configure read/write splitting for a master/standby instance, use the master and standby listener ports to connect to the master and standby nodes.

Figure 4-3 Public access addresses for a master/standby instance

Connection ⑦		
Password Protected	Yes	
Connection Address	redis-26ed6b50-c3f8-4a	dcs.com:6379 🗇 🖉
Read-only Address ③	redis-26ed6b50-c3f8	readonly.dcs.com:6379
IP Address	192.1 1:6379 🗇 🖉	
Public Access	On Off	
	Load balancer:elb-rs C	
	Listener: listener-master-26ed	3670445 🕐 Port: 6380
	Listener: listener-slave-26ed6b	70445 🗹 Port: 6381

• **Connection Address** and **IP Address** are the "domain name:port" and "IP address:port" for accessing Redis on a client within a VPC.

----End

# (Optional) Adding Private Load Balancer IP Addresses to the IP Address Whitelist of a Redis Instance

When the IP whitelist is enabled for Redis, add the private ELB IP addresses to the whitelist of the Redis instance to ensure that ELB can access the Redis instance.

1. Click the ELB address in **Public Access**.

Public Access	On Off		
	Load balancer;alb-r: i 🕐		

2. Copy **ID** of ELB.



- 3. Click **Private IPv4 address**.
- 4. On the **IP Addresses** tab page, in the second search box, filter private ELB IP addresses by resource ID (copied ELB ID).

IP Address	Resource ID $\Leftrightarrow$		Used By ⇔
192. 50	a3275	ic313767	Dedicated Load Balancer
192. 5	a3275	ic313767	Dedicated Load Balancer
192. 14	a3275i	ic313767	Dedicated Load Balancer
192. 33	a3275i	ic313767	Dedicated Load Balancer
90	a3275i	ic313767	Dedicated Load Balancer
192 35	a3275i	ic313767	Dedicated Load Balancer

5. Add all the private IP addresses of ELB to the IP whitelist of Redis. For details, see **Configuring DCS Redis Access Whitelist**.

# (Optional) Adding Public Client IP Addresses to the IP Address Whitelist of a Load Balancer

To configure a public IP address whitelist, add the public IP addresses of the DCS Redis instance to the IP address group of a load balancer. For details about how to add IP addresses to the group, see **IP Address Group**.

Only whitelisted public IP addresses can access the load balancer and then Redis. If no IP address whitelist is configured, all public IP addresses can access the load balancer by default.

# 4.2 Controlling DCS Redis Access

# 4.2.1 Configuring DCS Redis Access Whitelist

DCS helps you control access to your DCS instances in the following ways, depending on the deployment mode:

- To control access to DCS Redis 3.0, Memcached, and Redis 6.0 professional edition instances, you can use security groups. Whitelists are not supported. For details on how to configure a security group, see How Do I Configure a Security Group?
- To control access to DCS Redis 4.0 and later basic edition instances, you can use whitelists. Security groups are not supported.

The following describes how to manage whitelists of a Redis 4.0 or later basic edition instance to allow access only from whitelisted IP addresses. Enabling

whitelists only allows instance access from IP addresses within them, and only applies to new connections.

If no whitelists are added for the instance or the whitelist function is disabled, all IP addresses that can communicate with the VPC can access the instance.

# **Creating a Whitelist Group**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of a DCS instance.
- Step 5 Choose Instance Configuration > Whitelist. On the displayed page, click Create Whitelist Group.
- Step 6 In the Create Whitelist Group dialogue box, specify Group Name and IP Address/Range.

Table 4-1 Whitelist parameters

Parameter	Description	Example
Group Name	<ul> <li>Whitelist group name of the instance. A maximum of four whitelist groups can be created for each instance.</li> <li>Group naming rules:</li> <li>Start with a letter.</li> <li>4 to 64 characters.</li> <li>Only letters, digits, hyphens (-), and underscores (_) are allowed.</li> </ul>	DCS-test
IP Address/ Range	IP addresses or address segments of the instances allowed for access. A maximum of 100 IP addresses or IP address segments can be added to an instance. Use commas (,) to separate multiple IP addresses or address segments. Unsupported IP address and IP address range: 0.0.00 and 0.0.0.0/0.	10.10.10.1,10.10.10.10,19 2.168.0.0/16

### Step 7 Click OK.

The whitelist function takes effect immediately after the whitelist group is created. Only whitelisted IP addresses can access the instance. For persistent connections, the whitelist takes effect after reconnection.

## 

- To modify a whitelist: Click **Edit** on the whitelist page to modify the IP addresses or address segments of a whitelist.
- To delete a whitelist: Click **Delete** on the whitelist page to delete a whitelist group.
- To disable a whitelist: Click **Disable Whitelist** in the left corner of the whitelist tab page. After a whitelist is disabled, IP addresses within the same VPC as the instance can be used to access the instance. To enable the whitelist, click **Enable Whitelist**.

----End

# 4.2.2 Configuring a Redis Password

For security purposes, DCS provides password-protected instances. In addition, Redis can be accessed without a password. Use an instance access mode as required.

For a DCS instance that is used on the live network or contains important information, you are advised to set a password.

- To modify an instance password, see Changing an Instance Password.
- To change the access mode (password-protected or password-free), or to reset the password, see **Resetting an Instance Password**.

# **Suggestions for Password Security**

1. Hide the password when using redis-cli.

If the **-a** *<password>* option is used in redis-cli in Linux, the password is prone to leakage because it is logged and kept in the history. You are advised not to use **-a** *<password>* when running commands in redis-cli. After connecting to Redis, run the **auth** command to complete authentication. For example:

\$ redis-cli -h 192.168.0.148 -p 6379 redis 192.168.0.148:6379>**auth** *yourPassword* OK redis 192.168.0.148:6379>

2. Use interactive password authentication or switch between users with different permissions.

If the script involves DCS instance access, use interactive password authentication. To enable automatic script execution, manage the script as another user and authorize execution using sudo.

3. Use an encryption module in your application to encrypt the password.

# **Notes and Constraints**

- The instance must be in the **Running** state.
- Access the instance on a client using the latest password. Changing the password does not interrupt existing connections and the latest password is required upon a new connection.
- Only required for password-protected instances. For password-free ones, you can set a password by referring to **Resetting an Instance Password**.

• For security purposes, password-free access must be disabled when public access is enabled.

# **Changing an Instance Password**

### D NOTE

- After 5 consecutive incorrect password attempts, the account for accessing the chosen DCS instance will be locked for 5 minutes. Passwords cannot be changed during the lockout period. You can continue other operations.
- The new password takes effect immediately on the server without requiring a restart.
- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Choose **More** > **Change Password** in the row containing the chosen instance.
- **Step 5** The **Change Password** dialog box is displayed. Enter the old and new password, and confirm it.

The password must meet the following requirements:

- Cannot be left blank.
- Cannot be the same as the old password.
- Can be 8 to 64 characters long.
- Contain at least three of the following character types:
  - Lowercase letters
  - Uppercase letters
  - Digits
  - Special characters (`~!@#\$^&\*()-\_=+\|{},<.>/?)
- **Step 6** In the **Change Password** dialog box, click **OK** to confirm the password change.

----End

# **Resetting an Instance Password**

## D NOTE

- Disabling password protection may compromise security. You can set a password later by password resetting.
- For security purposes, password-free access must be disabled when public access is enabled.
- The system will display a success message only after the password is successfully reset on all nodes. If the reset fails, the instance will restart and the old password of the instance is still being used.
- Resetting passwords takes effect immediately without server restart.

### **Step 1** Log in to the **DCS console**.

**Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.

**Step 3** In the navigation pane, choose **Cache Manager**.

- **Step 4** To change the password setting for a DCS Redis instance, choose **More** > **Reset Password** in the **Operation** column in the row containing the chosen instance.
- **Step 5** In the **Reset Password** dialogue box, perform either of the following operations as required:
  - Change password-protected access to password-free access.

Switch the toggle for **Password-Free Access** and click **OK**.

• Change password-free access to password-protected access or reset the password.

Enter a password, confirm the password, and click **OK**.

The password must meet the following requirements:

- Cannot be left blank.
- Cannot be the same as the old password.
- Can be 8 to 64 characters long.
- Contain at least three of the following character types:
  - Lowercase letters
  - Uppercase letters
  - Digits
  - Special characters (`~!@#\$^&\*()-\_=+\|{},<.>/?)

----End

# 4.2.3 Transmitting DCS Redis Data with Encryption Using SSL

**Single-node, master/standby, and Redis Cluster basic edition DCS Redis 6.0/7.0** instances support SSL encryption to ensure data transmission security. This function is not available for other instance versions. RESP (Redis Serialization Protocol), the communication protocol of Redis, only supports plaintext transmission in versions earlier than Redis 6.0.

# Notes and Constraints

- Due to SSL encryption, SSL and client IP pass-through cannot be enabled at the same time. Encrypted links do not carry client IPs.
- Enabling SSL will deteriorate read/write performance.
- Enabling or disabling SSL will restart the instance and disconnect it for a few seconds. Wait until off-peak hours and ensure that your application can reconnect.
- The restart cannot be undone. For single-node DCS instances and other instances where AOF persistence is disabled (**appendonly** is set to **no**), data will be cleared and ongoing backup tasks will be stopped. Exercise caution when performing this operation.

# **Enabling or Disabling SSL**

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

- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- Step 3 In the navigation pane, choose Cache Manager.
- **Step 4** On the **Cache Manager** page, click a DCS instance.
- Step 5 In the navigation pane, choose SSL.
- **Step 6** Click A next to **SSL Certificate** to enable or disable SSL.

# NOTICE

- Enabling or disabling SSL will restart the instance and disconnect it for a few seconds. Wait until off-peak hours and ensure that your application can reconnect.
- The restart cannot be undone. For single-node DCS instances and other instances where AOF persistence is disabled (**appendonly** is set to **no**), data will be cleared and ongoing backup tasks will be stopped. Exercise caution when performing this operation.
- Step 7 Click Download Certificate to download the SSL certificate.
- **Step 8** Decompress the SSL certificate and upload the decompressed **ca.crt** file to the server where the Redis client is located.
- **Step 9** Add the path of the **ca.crt** file to the command for connecting to the instance. For example, to access an instance on redis-cli, see **Connecting to Redis on redis-cli**.

----End

# 4.2.4 Configuring DCS Redis ACL Users

If you need multiple accounts for a Redis instance, use ACL to create users. ACL users support read-only or read/write permissions.

# **Notes and Constraints**

- Currently, DCS Redis 4.0 and 5.0 instances support user management by default. To use this function for DCS Redis 6.0 instances, contact customer service to enable the whitelist and upgrade the instance engine.
- This function can only be enabled for DCS Redis professional instances with a password and is not supported for DCS Redis professional instances with **Password Protected** disabled.
- Up to 18 users can be created for each instance.

# **Configuring DCS Redis ACL Users**

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

- **Step 2** Click <sup>(2)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click an instance.
- **Step 5** Choose **User Management** in the navigation pane.

The user whose username is **default** is the instance's default user. The default user has read and write permissions and their password is the instance's password.

Step 6 Click Create User.

A password-free Redis instance can only be accessed by the default user. Normal users are invalid. To use a normal user, click **Reset Password** in the row containing the default user to disable **Password Protected** for the default user.

- **Step 7** Specify the **Username** and **Description**. Select **Read-only** or **Read/Write**. Specify the **Password** and confirm it.
- Step 8 Click OK.

### NOTICE

A normal ACL user connects to an instance with password *{username:password}*.

• When **using redis-cli** to connect to an instance, the default user runs the following command:

./redis-cli -h {dcs\_instance\_address} -p 6379 -a {password}

• A normal ACL user runs the following command:

./redis-cli -h {dcs\_instance\_address} -p 6379 -a {username:password}

#### Figure 4-4 User management

User Management					
Create User Delete					
Q Select a property or enter a keyword.					(a)
Username 👌	Type 😂	Status 🖯	Permissions 🕀	Description \ominus	Operation
readonly	Normal	Available	Read-only	-	Change Password Reset Password More
default	Default	Available	Read/Write	-	Change Password Reset Password

----End

# **More Operations**

The following operations can be performed on normal users.

Table 4-2 Operation

Operation	Description
Changing a password	Locate the row that contains the desired normal user and click <b>Change Password</b> in the <b>Operation</b> column.

Description
If password is forgot, locate the row that contains the normal user and click <b>Reset Password</b> in the <b>Operation</b> column.
Locate the row that contains the normal user. Choose <b>More</b> > <b>Modify Permission</b> in the <b>Operation</b> column. The <b>Read-only</b> or <b>Read/Write</b> permissions can be granted.
Locate the row that contains the normal user. Choose <b>More</b> > <b>Edit Description</b> in the <b>Operation</b> column.
Locate the row that contains the normal user. Choose <b>More</b> > <b>Delete</b> in the <b>Operation</b> column.
Select the normal users to be deleted and click <b>Delete</b> above the list.

# 4.3 Connecting to Redis on a Client

# 4.3.1 Connecting to Redis on redis-cli

This section describes how to access a DCS Redis instance on redis-cli. For more information about how to use other Redis clients, visit **the Redis official website**.

The following operations are based on an example of accessing a Redis instance on a client on an elastic cloud server (ECS).

To access a Redis 3.0 instance over a public network, see **Connecting to Redis 3.0** over a Public Network on redis-cli.

# Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- An ECS has been created. For details about how to create an ECS, see Purchasing a Custom ECS.
- If the ECS runs the Linux OS, ensure that the GCC compilation environment has been installed on the ECS.

Run the following command to install GCC on the ECS if needed:

```
yum install -y make
yum install -y pcre-devel
yum install -y zlib-devel
yum install -y libevent-devel
yum install -y openssl-devel
yum install -y gcc-c++
```

 The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

# Connecting to Redis on redis-cli (Linux)

**Step 1** View the IP address/domain name and port of the DCS Redis instance to be accessed.

See Viewing Instance Details to obtain Connection Address or IP Address in the instance information. For more information, see Should I Use a Domain Name or an IP Address to Connect to a DCS Redis Instance?.

For public access to Redis on a client, see **Enabling Public Access to Redis and Obtaining the Access Addresses** to obtain the instance public addresses and ports.



Connection ⑦	
Password Protected	No
Connection Address	redis-1995-1993-1994-1993-1999-1993 red Di-north-4.dcs.myhuaweicloud.com:6379 🗇 🖉
IP Address	10.0.0.70:6379 🗇
Public Access	View Documentation

# **NOTE**

- The following uses the port 6379. Replace 6379 with the actual port.
- For Proxy Cluster Redis 4.0 and later instances, **Connection Address** and **IP Address** are load balancer addresses. The system distributes requests to different proxies.
- You can use **Backend Addresses** to directly connect to the specified proxy node of a Proxy Cluster DCS Redis 3.0 instance.

#### **Step 2** Install redis-cli.

The following steps assume that your client is installed on the Linux OS.

- 1. Log in to the ECS.
- Run the following command to download the source code package of your Redis client from https://download.redis.io/releases/redis-6.2.13.tar.gz: wget http://download.redis.io/releases/redis-6.2.13.tar.gz

#### **NOTE**

The following uses redis-6.2.13 as an example. For details, see the **Redis official website**.

- Run the following command to decompress the source code package of your Redis client: tar -xzf redis-6.2.13.tar.gz
- 4. Run the following commands to go to the Redis directory and compile the source code of your Redis client: cd redis-6.2.13 make cd src

## D NOTE

If the source code of your Redis client is v6.0 and later, and redis-cli that supports TLS/SSL is required, replace the **make** command with **make BUILD\_TLS=yes** to enable TLS.

#### **Step 3** Access the DCS Redis instance.

### • Access a DCS instance of a type other than Redis Cluster.

Perform the following procedure to access a single-node, master/standby, read/write splitting, or Proxy Cluster instance.

a. Run the following command to access the chosen DCS Redis instance: ./redis-cli -h {*dcs\_instance\_address*} -p {*port*}

**{dcs\_instance\_address}** indicates the IP address/domain name of the DCS instance and **{port}** is the port used for accessing the instance. The IP address/domain name and port number are obtained in **Step 1**.

The following example uses the domain name address of a DCS Redis instance. Change the connection address and port as required.

[root@ecs-redis src]# ./redis-cli -h redis-069949a-dcs-lxy.dcs.huaweicloud.com -p 6379 redis-069949a-dcs-lxy.dcs.huaweicloud.com:6379>

b. If you have set a password for the DCS instance, enter the password in this step. You can read and write cached data only after the password is verified. Skip this step if the instance is not password-protected. auth {password}

**{password}** indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

The command output is as follows:

redis-069949a-dcs-lxy.dcs.huaweicloud.com:6379> auth \*\*\*\*\*\*\*

OK redis-069949a-dcs-lxy.dcs.huaweicloud.com:6379>

### • Access a DCS instance of the Redis Cluster type.

Perform the following procedure to access a Redis Cluster instance.

a. Run the following command to access the chosen DCS Redis instance: ./redis-cli -h {*dcs\_instance\_address*} -p {*port*} -a {*password*} -c

**{dcs\_instance\_address}** indicates the IP address/domain name of the DCS Redis instance, **{port}** is the port used for accessing the instance, **{password}** is the password of the instance, and **-c** is used for accessing Redis Cluster nodes. The IP address/domain name and port number are obtained in **Step 1**.

## D NOTE

- When connecting to a Redis Cluster instance using redis-cli, ensure that c is added to the command. Otherwise, the connection will fail.
- To connect to Redis on a client over a private network, you can set {dcs\_instance\_address} to Connection Address or IP Address in the Connection section, or IP Address in the Instance Topology section. The addresses can be obtained on the instance basic information page on the console, as shown in Figure 4-6.
- For a password-protected instance, you do not need to enter the instance access password -a {password}. If you have forgotten the password or need to reset the password, see Resetting an Instance Password.
- The IP Address field provides multiple IP addresses. You can use any of them to connect to the instance. The CRC16(key) mod 16384 algorithm is used to compute what is the hash slot of a given key. For higher reliability, configure all IP addresses.
- By using the IP Address in the Instance Topology section, you can connect to the specified shard.

**Figure 4-6** Obtaining the addresses for connecting to a Redis Cluster DCS instance



- The following example uses the IP address of a DCS Redis instance. Change the IP address and port as required. root@ecs-redis:~/redis-6.2.13/src# ./redis-cli -h 192.168.0.85 -p 6379 -a \*\*\*\*\*\* -c 192.168.0.85:6379>
- The following example uses the domain name of a DCS Redis instance. Change the domain name and port as required. root@ecs-redis:~/redis-6.2.13/src# ./redis-cli -h redis-51e463c-dcs-lxy.dcs.huaweicloud.com -p 6379 -a \*\*\*\*\*\* -c redis-51e463c-dcs-lxy.dcs.huaweicloud.com:6379>
- b. Run the **cluster nodes** command to view the Redis Cluster node information:

Each shard in a Redis Cluster has a master and a replica by default. The proceeding command provides all the information of cluster nodes.
0988ae8fd3686074c9afdcce73d7878c81a33ddc 192.168.0.231:6379@16379 slave f0141816260ca5029c56333095f015c7a058f113 0 1568084030 000 3 connected 1a32d809c0b743bd83b5e1c277d5d201d0140b75 192.168.0.85:6379@16379 myself,master - 0 1568084030000 2 connected 5461-10922 c8ad7af9a12cce3c8e416fb67bd6ec9207f0082d 192.168.0.130:6379@16379 slave 1a32d809c0b743bd83b5e1c277d5d201d0140b75 0 1568084031 000 2 connected 7ca218299c254b5da939f8e60a940ac8171adc27 192.168.0.22:6379@16379 master - 0 1568084030000 1 connected 0-5460 f0141816260ca5029c56333095f015c7a058f113 192.168.0.170:6379@16379 master - 0 1568084031992 3 connected 10923-16383 19b1a400815396c6223963b013ec934a657bdc52 192.168.0.161:6379@16379 slave 7ca218299c254b5da939f8e60a940ac8171adc27 0 1568084031 000 1 connected Write operations can only be performed on master nodes. The CRC16(key) mod 16384 algorithm is used to compute what is the hash slot of a given key. As shown in the following, the value of CRC16 (KEY) mode 16384 determines the hash slot that a given key is located at and redirects the client to the node where the hash slot is located at. 192.168.0.170:6379> set hello world -> Redirected to slot [866] located at 192.168.0.22:6379 OK 192.168.0.22:6379> set happy day OK 192.168.0.22:6379> set abc 123 -> Redirected to slot [7638] located at 192.168.0.85:6379 OK 192.168.0.85:6379> get hello

-> Redirected to slot [866] located at 192.168.0.22:6379 "world" 192.168.0.22:6379> get abc -> Redirected to slot [7638] located at 192.168.0.85:6379 "123" 192.168.0.85:6379>

192.168.0.85:6379> cluster nodes

#### **NOTE**

If **SSL** is enabled for a DCS Redis 6.0/7.0 basic edition instance, configure an SSL certificate path.

- Run the following command to connect to a Redis Cluster instance: ./redis-cli -h {dcs\_instance\_address} -p {port} -a {password} -c --tls --cacert {certification file path}
- Run the following command to connect to a single-node or master/standby instance: ./redis-cli -h {dcs\_instance\_address} -p {port} -a {password} --tls --cacert {certification file path}

To connect to Redis with SSL encryption, use redis-cli 6.x or later.

----End

# Connecting to Redis on redis-cli (Windows)

Click **here** to download the Redis client installation package for Windows. Decompress the package in any directory, open the CLI tool **cmd.exe**, and go to the directory. Then, run the following command to access the DCS Redis instance:

redis-cli.exe -h XXX -p 6379

**XXX** indicates the IP address/domain name of the DCS instance and **6379** is an example port number used for accessing the DCS instance. For details about how

to obtain the IP address/domain name and port number, see **Viewing Instance Details**.

# 4.3.2 Connecting to Redis on Jedis (Java)

This section describes how to access a Redis instance on Jedis. For more information about how to use other Redis clients, visit **the Redis official website**.

Spring Data Redis is already integrated with **Jedis** and **Lettuce** for Spring Boot projects. Spring Boot 1.x is integrated with Jedis, and Spring Boot 2.x is integrated with Lettuce. To use Jedis in Spring Boot 2.x and later, you need to solve Lettuce dependency conflicts.

#### **Notes and Constraints**

Springboot 2.3.12.RELEASE or later is required. Jedis **3.10.0** or later is required.

To access a Redis 7.0 instance, use a Jedis **5.0.0** or later client. **5.1.1** and later versions are recommended.

#### Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- View the IP address/domain name and port of the DCS Redis instance to be accessed. For details, see Viewing and Modifying Basic Settings of a DCS Instance.
- The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

## **Pom Configuration**

<!-- import spring-data-redis --> <dependency> <groupId>org.springframework.boot</groupId> <artifactId>spring-boot-starter-data-redis</artifactId> <!--In Spring Boot 2.0, Lettuce is used by default. To use Jedis, solve dependency conflicts.--> <exclusions> <exclusion> <groupId>io.lettuce</groupId> <artifactId>lettuce-core</artifactId> </exclusion> </exclusions> </dependency> <!--Jedis dependency> <dependency> <groupId>redis.clients</groupId> <artifactId>jedis</artifactId> <version>\${jedis.version}<version> </dependency>

## application.properties Configuration

 Single-node, master/standby, read/write splitting, and Proxy Cluster #Redis host spring.redis.host=<host> #Redis port spring.redis.port=<port> #Redis database number spring.redis.database=0 #Redis password spring.redis.password=<password> #Redis read/write timeout spring.redis.timeout=2000 #Whether to enable connection pooling spring.redis.jedis.pool.enabled=true #Minimum connections in the pool spring.redis.jedis.pool.min-idle=50 #Maximum idle connections in the pool spring.redis.jedis.pool.max-idle=200 #Maximum connections in the pool spring.redis.jedis.pool.max-active=200 #Maximum amount of time a connection allocation should block before throwing an exception when the pool is exhausted. The default value -1 indicates to wait indefinitely. spring.redis.jedis.pool.max-wait=3000 #Interval for checking and evicting idle connection. Default: 60s. spring.redis.jedis.pool.time-between-eviction-runs=60S **Redis Cluster** #Redis Cluster node connection information

spring.redis.cluster.nodes=*<ip:port>*,*<ip:port>*,*<ip:port>*, #Redis Cluster password spring.redis.password=<password> #Redis Cluster max. redirecting times spring.redis.cluster.max-redirects=3 #Redis read/write timeout spring.redis.timeout=2000 #Whether to enable connection pooling spring.redis.jedis.pool.enabled=true #Minimum connections in the pool spring.redis.jedis.pool.min-idle=50 #Maximum idle connections in the pool spring.redis.jedis.pool.max-idle=200 #Maximum connections in the pool spring.redis.jedis.pool.max-active=200 #Maximum amount of time a connection allocation should block before throwing an exception when the pool is exhausted. The default value -1 indicates to wait indefinitely. spring.redis.jedis.pool.max-wait=3000 #Interval for checking and evicting idle connections. Default: 60s. spring.redis.jedis.pool.time-between-eviction-runs=60S

## **Bean Configuration**

• Single-node, master/standby, read/write splitting, and Proxy Cluster import java.time.Duration;

import org.springframework.beans.factory.annotation.Value; import org.springframework.context.annotation.Bean; import org.springframework.context.annotation.Configuration; import org.springframework.data.redis.connection.RedisConnectionFactory; import org.springframework.data.redis.connection.RedisStandaloneConfiguration; import org.springframework.data.redis.connection.jedis.JedisClientConfiguration; import org.springframework.data.redis.connection.jedis.JedisClientConfiguration; import org.springframework.data.redis.connection.jedis.JedisClientConfiguration;

import redis.clients.jedis.JedisPoolConfig;

@Configuration
public class RedisConfiguration {

@Value("\${redis.host}")
private String redisHost;

@Value("\${redis.port:6379}")
private Integer redisPort = 6379;

@Value("\${redis.database:0}")
private Integer redisDatabase = 0;

```
@Value("${redis.password:}")
private String redisPassword;
```

@Value("\${redis.connect.timeout:3000}") private Integer redisConnectTimeout = 3000;

@Value("\${redis.read.timeout:2000}") private Integer redisReadTimeout = 2000;

@Value("\${redis.pool.minSize:50}") private Integer redisPoolMinSize = 50;

@Value("\${redis.pool.maxSize:200}") private Integer redisPoolMaxSize = 200;

@Value("\${redis.pool.maxWaitMillis:3000}") private Integer redisPoolMaxWaitMillis = 3000;

@Value("\${redis.pool.softMinEvictableIdleTimeMillis:1800000}") private Integer redisPoolSoftMinEvictableIdleTimeMillis = 30 \* 60 \* 1000;

@Value("\${redis.pool.timeBetweenEvictionRunsMillis:60000}") private Integer redisPoolBetweenEvictionRunsMillis = 60 \* 1000;

@Bean

public RedisConnectionFactory redisConnectionFactory(JedisClientConfiguration clientConfiguration) {

RedisStandaloneConfiguration standaloneConfiguration = new RedisStandaloneConfiguration(); standaloneConfiguration.setHostName(redisHost); standaloneConfiguration.setPort(redisPort); standaloneConfiguration.setDatabase(redisDatabase); standaloneConfiguration.setPassword(redisPassword);

return new JedisConnectionFactory(standaloneConfiguration, clientConfiguration);

}

@Bean

public JedisClientConfiguration clientConfiguration() {

JedisClientConfiguration clientConfiguration = JedisClientConfiguration.builder() .connectTimeout(Duration.ofMillis(redisConnectTimeout)) .readTimeout(Duration.ofMillis(redisReadTimeout)) .usePooling().poolConfig(redisPoolConfig()) .build();

return clientConfiguration;

}

private JedisPoolConfig redisPoolConfig() {

JedisPoolConfig poolConfig = new JedisPoolConfig(); //Minimum connections in the pool poolConfig.setMinIdle(redisPoolMinSize); //Maximum idle connections in the pool poolConfig.setMaxIdle(redisPoolMaxSize); //Maximum total connections in the pool poolConfig.setMaxTotal(redisPoolMaxSize); //Wait when pool is exhausted? Set to true to wait. To validate setMaxWait, it has to be true. poolConfig.setBlockWhenExhausted(true); //Longest time to wait for connection after pool is exhausted. The default value -1 indicates to wait indefinitely. poolConfig.setMaxWaitMillis(redisPoolMaxWaitMillis); //Set to true to enable connectivity test on creating connections. Default: false. poolConfig.setTestOnCreate(false); //Set to true to enable connectivity test on borrowing connections. Default: false. Set to false for heavy-traffic services to reduce overhead. poolConfig.setTestOnBorrow(true);

//Set to true to enable connectivity test on returning connections. Default: false. Set to false for

heavy-traffic services to reduce overhead. poolConfig.setTestOnReturn(false); //Indicates whether to check for idle connections. If this is set to false, idle connections are not evicted. poolConfig.setTestWhileIdle(true); //Duration after which idle connections are evicted. If the idle duration is greater than this value and the maximum number of idle connections is reached, idle connections are directly evicted. poolConfig.setSoftMinEvictableIdleTimeMillis(redisPoolSoftMinEvictableIdleTimeMillis); //Disable MinEvictableIdleTimeMillis(). poolConfig.setMinEvictableIdleTimeMillis(-1); //Interval for checking and evicting idle connections. Default: 60s. poolConfig.setTimeBetweenEvictionRunsMillis(redisPoolBetweenEvictionRunsMillis); return poolConfig; } } **Redis Cluster** import java.time.Duration; import java.util.ArrayList; import java.util.List; import org.springframework.beans.factory.annotation.Value; import org.springframework.context.annotation.Bean; import org.springframework.context.annotation.Configuration; import org.springframework.data.redis.connection.RedisClusterConfiguration; import org.springframework.data.redis.connection.RedisConnectionFactory; import org.springframework.data.redis.connection.RedisNode; import org.springframework.data.redis.connection.jedis.JedisClientConfiguration; import org.springframework.data.redis.connection.jedis.JedisConnectionFactory; import redis.clients.jedis.JedisPoolConfig; @Configuration public class RedisConfiguration { @Value("\${redis.cluster.nodes}") private String redisClusterNodes; @Value("\${redis.password:}") private String redisPassword; @Value("\${redis.connect.timeout:3000}") private Integer redisConnectTimeout = 3000; @Value("\${redis.read.timeout:2000}") private Integer redisReadTimeout = 2000; @Value("\${redis.pool.minSize:50}") private Integer redisPoolMinSize = 50; @Value("\${redis.pool.maxSize:200}") private Integer redisPoolMaxSize = 200; @Value("\${redis.pool.maxWaitMillis:3000}") private Integer redisPoolMaxWaitMillis = 3000; @Value("\${redis.pool.softMinEvictableIdleTimeMillis:1800000}") private Integer redisPoolSoftMinEvictableIdleTimeMillis = 30 \* 60 \* 1000; @Value("\${redis.pool.timeBetweenEvictionRunsMillis:60000}") private Integer redisPoolBetweenEvictionRunsMillis = 60 \* 1000; @Bean public RedisConnectionFactory redisConnectionFactory(JedisClientConfiguration clientConfiguration) { RedisClusterConfiguration clusterConfiguration = new RedisClusterConfiguration(); List<RedisNode> clusterNodes = new ArrayList<>();

for (String clusterNodeStr : redisClusterNodes.split(",")) {

```
String[] nodeInfo = clusterNodeStr.split(":");
        clusterNodes.add(new RedisNode(nodeInfo[0], Integer.valueOf(nodeInfo[1])));
     }
     clusterConfiguration.setClusterNodes(clusterNodes);
     clusterConfiguration.setPassword(redisPassword);
     clusterConfiguration.setMaxRedirects(3);
     return new JedisConnectionFactory(clusterConfiguration, clientConfiguration);
  }
  @Bean
  public JedisClientConfiguration clientConfiguration() {
     JedisClientConfiguration clientConfiguration = JedisClientConfiguration.builder()
           .connectTimeout(Duration.ofMillis(redisConnectTimeout))
           .readTimeout(Duration.ofMillis(redisReadTimeout))
           .usePooling().poolConfig(redisPoolConfig())
           .build();
     return clientConfiguration;
  }
  private JedisPoolConfig redisPoolConfig() {
     JedisPoolConfig poolConfig = new JedisPoolConfig();
     //Minimum connections in the pool
     poolConfig.setMinIdle(redisPoolMinSize);
     //Maximum idle connections in the pool
     poolConfig.setMaxIdle(redisPoolMaxSize);
     //Maximum total connections in the pool
     poolConfig.setMaxTotal(redisPoolMaxSize);
     //Wait when pool is exhausted? Set to true to wait. To validate setMaxWait, it has to be true.
     poolConfig.setBlockWhenExhausted(true);
     //Longest time to wait for connection after pool is exhausted. The default value -1 indicates to
wait indefinitely.
     poolConfig.setMaxWaitMillis(redisPoolMaxWaitMillis);
     //Set to true to enable connectivity test on creating connections. Default: false.
     poolConfig.setTestOnCreate(false);
     //Set to true to enable connectivity test on borrowing connections. Default: false. Set to false for
heavy-traffic services to reduce overhead.
     poolConfig.setTestOnBorrow(true);
     //Set to true to enable connectivity test on returning connections. Default: false. Set to false for
heavy-traffic services to reduce overhead.
     poolConfig.setTestOnReturn(false);
     //Indicates whether to check for idle connections. If this is set to false, idle connections are not
evicted.
     poolConfig.setTestWhileIdle(true);
     //Duration after which idle connections are evicted. If the idle duration is greater than this value
and the maximum number of idle connections is reached, idle connections are directly evicted.
     poolConfig.setSoftMinEvictableIdleTimeMillis(redisPoolSoftMinEvictableIdleTimeMillis);
     //Disable MinEvictableIdleTimeMillis().
     poolConfig.setMinEvictableIdleTimeMillis(-1);
     //Interval for checking and evicting idle connections. Default: 60s.
     poolConfig.setTimeBetweenEvictionRunsMillis(redisPoolBetweenEvictionRunsMillis);
     return poolConfig;
  }
}
```

## (Optional) Configuring SSL Connections

If SSL is enabled for the instance, use the following content to replace the JedisClientConfiguration construction method clientConfiguration() in **Bean Configuration** for connecting to the instance with SSL. For details about whether your DCS Redis instances support SSL, see **Transmitting DCS Redis Data with Encryption Using SSL**.

```
@Bean
public JedisClientConfiguration clientConfiguration() throws Exception {
  JedisClientConfiguration.JedisClientConfigurationBuilder configurationBuilder
     = JedisClientConfiguration.builder()
     .connectTimeout(Duration.ofMillis(redisConnectTimeout))
     .readTimeout(Duration.ofMillis(redisReadTimeout));
  configurationBuilder.usePooling().poolConfig(redisPoolConfig());
  configurationBuilder.useSsl().sslSocketFactory(getTrustStoreSslSocketFactory());
  return configurationBuilder.build();
private SSLSocketFactory getTrustStoreSslSocketFactory() throws Exception{
  //Load the CA certificate in the user-defined path based on service requirements.
  CertificateFactory cf = CertificateFactory.getInstance("X.509");
  Certificate ca;
  try (InputStream is = new FileInputStream("./ca.crt")) {
     ca = cf.generateCertificate(is);
  }
  //Create keystore.
  String keyStoreType = KeyStore.getDefaultType();
  KeyStore keyStore = KeyStore.getInstance(keyStoreType);
  keyStore.load(null, null);
  keyStore.setCertificateEntry("ca", ca);
  //Create TrustManager.
  TrustManagerFactory trustManagerFactory = TrustManagerFactory.getInstance(
     TrustManagerFactory.getDefaultAlgorithm());
  trustManagerFactory.init(keyStore);
  //Create SSLContext.
  SSLContext context = SSLContext.getInstance("TLS");
  context.init(null, trustManagerFactory.getTrustManagers(), new SecureRandom());
  return context.getSocketFactory();
```

## **Parameter Description**

Parameter	Default Value	Description
hostName	localhost	IP address/domain name for connecting to a DCS Redis instance
port	6379	Port number
database	0	Database number. Default: 0.
password	-	Redis instance password

 Table 4-3
 RedisStandaloneConfiguration parameters

Table 4-4 RedisClusterConfiguration	parameters
-------------------------------------	------------

Parameter	Description
clusterNodes	Cluster node connection information, including the node IP address and port number
maxRedirects	Maximum redirecting times

Parameter	Description	
password	Password	

## Table 4-5 JedisPoolConfig parameters

Parameter	Default Value	Description	
minIdle	-	Minimum connections in the connection pool	
maxIdle	-	Maximum idle connections in the connection pool	
maxTotal	-	Maximum total connections in the connection pool	
blockWhenExha usted	true	Indicates whether to wait after the connection pool is exhausted. <b>true</b> : Wait. <b>false</b> : Do not wait. To validate <b>maxWaitMillis</b> , this parameter must be set to <b>true</b> .	
maxWaitMillis	-1	Maximum amount of time (in milliseconds) to wait for connection after the connection pool is exhausted. The default value <b>-1</b> indicates to wait indefinitely.	
testOnCreate	false	Indicates whether to enable connectivity test on creating connections. <b>false</b> : Disable. <b>true</b> : Enable.	
testOnBorrow	false	Indicates whether to enable connectivity test on obtaining connections. <b>false</b> : Disable. <b>true</b> : Enable. For heavy-traffic services, set this parameter to <b>false</b> to reduce overhead.	
testOnReturn	false	Indicates whether to enable connectivity test on returning connections. <b>false</b> : Disable. <b>true</b> : Enable. For heavy-traffic services, set this parameter to <b>false</b> to reduce overhead.	
testWhileIdle	false	Indicates whether to check for idle connections. If this parameter is set to <b>false</b> , idle connections are not evicted. Recommended value: <b>true</b> .	
softMinEvictabl eIdleTimeMillis	1800000	Duration (in milliseconds) after which idle connections are evicted. If the idle duration is greater than this value and the maximum number of idle connections is reached, idle connections are directly evicted.	

Parameter	Default Value	Description
minEvictableIdle TimeMillis	60000	Minimum amount of time (in milliseconds) a connection may remain idle in the pool before it is eligible for eviction. The recommended value is <b>-1</b> , indicating that <b>softMinEvictableI-dleTimeMillis</b> is used instead.
timeBetweenEvi ctionRunsMillis	60000	Interval (in milliseconds) for checking and evicting idle connections.

 Table 4-6 JedisClientConfiguration parameters

Parameter	Default Value	Description
connectTimeout	2000	Connection timeout interval, in milliseconds.
readTimeout	2000	Timeout interval for waiting for a response, in milliseconds.
poolConfig	-	Pool configurations. For details, see JedisPoolConfig.

# **Suggestion for Configuring DCS Instances**

• Connection pool configuration

**NOTE** 

The following calculation is applicable only to common service scenarios. You can customize it based on your service requirements.

There is no standard connection pool size. You can configure one based on your service traffic. The following formulas are for reference:

- Minimum number of connections = (QPS of a single node accessing Redis)/(1000 ms/Average time spent on a single command)
- Maximum number of connections = (QPS of a single node accessing Redis)/(1000 ms/Average time spent on a single command) x 150%

For example, if the QPS of a service application is about 10,000, each request needs to access Redis 10 times (that is, 100,000 accesses to Redis every second), and the service application has 10 hosts, the calculation is as follows:

QPS of a single node accessing Redis = 100,000/10 = 10,000

Average time spent on a single command = 20 ms (Redis takes 5 ms to 10 ms to process a single command under normal conditions. If network jitter occurs, it takes 15 ms to 20 ms.)

Minimum number of connections = 10,000/(1000 ms/20 ms) = 200

Maximum number of connections =  $10,000/(1000 \text{ ms}/20 \text{ ms}) \times 150\% = 300$ 

# 4.3.3 Connecting to Redis on Lettuce (Java)

This section describes how to access a Redis instance on Lettuce. For more information about how to use other Redis clients, visit **the Redis official website**.

Spring Data Redis is already integrated with **Jedis** and **Lettuce** for Spring Boot projects. In addition, Spring Boot 1.x is integrated with Jedis, and Spring Boot 2.x with Lettuce. Therefore, you do not need to import Lettuce in Spring Boot 2.x and later projects.

#### **Notes and Constraints**

Springboot 2.3.12.RELEASE or later is required. Lettuce **6.3.0.RELEASE** or later is required. Netty 4.1.100.Final or later is required.

#### **Prerequisites**

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- View the IP address/domain name and port of the DCS Redis instance to be accessed. For details, see Viewing and Modifying Basic Settings of a DCS Instance.
- The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

#### **Pom Configuration**

<!-- Enable Spring Data Redis. Lettuce-supported SDK is integrated by default. --> <dependency> <groupId>org.springframework.boot</groupId> <artifactId>spring-boot-starter-data-redis</artifactId> </dependency> <dependency> <groupId>io.lettuce</groupId> <artifactId>lettuce-core</artifactId> <version>6.3.0.RELEASE</version> </dependency> <dependency> <groupId>io.netty</groupId> <artifactId>netty-transport-native-epoll</artifactId> <version>4.1.100.Final</version> <classifier>linux-x86\_64</classifier> </dependency>

## application.properties Configuration

Single-node, master/standby, read/write splitting, and Proxy Cluster

#Redis host
spring.redis.host=<host>
#Redis port
spring.redis.port=<port>
#Redis database number
spring.redis.database=0
#Redis password
spring.redis.password=<password>
#Redis read/write timeout
spring.redis.timeout=2000

#### **Redis Cluster**

#Redis Cluster node information spring.redis.cluster.nodes=<ip:port>,<ip:port>,<ip:port> #Redis Cluster max redirecting times spring.redis.cluster.max-redirects=3 #Redis Cluster node password spring.redis.password=<password> #Redis Cluster timeout spring.redis.timeout=2000 #Enable adaptive topology refresh spring.redis.lettuce.cluster.refresh.adaptive=true #Enable topology refresh every 10 seconds spring.redis.lettuce.cluster.refresh.period=10S

#### **Bean Configuration**

Single-node, master/standby, read/write splitting, and Proxy Cluster import java.time.Duration;

import org.springframework.beans.factory.annotation.Value; import org.springframework.context.annotation.Bean; import org.springframework.context.annotation.Configuration; import org.springframework.data.redis.connection.RedisConnectionFactory; import org.springframework.data.redis.connection.RedisStandaloneConfiguration; import org.springframework.data.redis.connection.lettuce.LettuceClientConfiguration; import org.springframework.data.redis.connection.lettuce.LettuceConnectionFactory;

import io.lettuce.core.ClientOptions; import io.lettuce.core.SocketOptions;

\* Lettuce non-pooling configuration (use either this or the application.properties configuration) @Configuration public class RedisConfiguration {

@Value("\${redis.host}") private String redisHost;

@Value("\${redis.port:6379}") private Integer redisPort = 6379;

@Value("\${redis.database:0}") private Integer redisDatabase = 0;

@Value("\${redis.password:}") private String redisPassword;

@Value("\${redis.connect.timeout:2000}") private Integer redisConnectTimeout = 2000;

@Value("\${redis.read.timeout:2000}") private Integer redisReadTimeout = 2000;

\* TCP\_KEEPALIVE configuration parameters:

- \* A keepalive interval = TCP\_KEEPALIVE\_TIME = 30
- \* Idle duration before keepalive = TCP\_KEEPALIVE\_TIME/3 = 10
- keepalive xx times before disconnect = TCP\_KEEPALIVE\_COUNT = 3

private static final int TCP\_KEEPALIVE\_TIME = 30;

\* TCP\_USER\_TIMEOUT Idle duration limit, to address Lettuce timeout. \* refer: https://github.com/lettuce-io/lettuce-core/issues/2082

private static final int TCP\_USER\_TIMEOUT = 30;

@Bean

public RedisConnectionFactory redisConnectionFactory(LettuceClientConfiguration

clientConfiguration) {

```
RedisStandaloneConfiguration standaloneConfiguration = new RedisStandaloneConfiguration();
     standaloneConfiguration.setHostName(redisHost);
     standaloneConfiguration.setPort(redisPort);
     standaloneConfiguration.setDatabase(redisDatabase);
     standaloneConfiguration.setPassword(redisPassword);
     LettuceConnectionFactory connectionFactory = new
LettuceConnectionFactory(standaloneConfiguration, clientConfiguration);
     connectionFactory.setDatabase(redisDatabase);
     return connectionFactory;
  }
  @Bean
  public LettuceClientConfiguration clientConfiguration() {
     SocketOptions socketOptions = SocketOptions.builder()
       .keepAlive(SocketOptions.KeepAliveOptions.builder()
          // A keepalive interval
          .idle(Duration.ofSeconds(TCP_KEEPALIVE_TIME))
          // Idle duration before keepalive
          .interval(Duration.ofSeconds(TCP_KEEPALIVE_TIME/3))
          // keepalive xx times before disconnect
          .count(3)
          // Whether to keep connections alive.
          .enable()
          .build())
       .tcpUserTimeout(SocketOptions.TcpUserTimeoutOptions.builder()
          // Addressing timeouts caused by RST on the server
          .tcpUserTimeout(Duration.ofSeconds(TCP_USER_TIMEOUT))
          .enable()
          .build())
       // TCP connection timeout setting
       .connectTimeout(Duration.ofMillis(redisConnectTimeout))
       .build();
     ClientOptions clientOptions = ClientOptions.builder()
          .autoReconnect(true)
          .pingBeforeActivateConnection(true)
          .cancelCommandsOnReconnectFailure(false)
          .disconnectedBehavior(ClientOptions.DisconnectedBehavior.ACCEPT_COMMANDS)
          .socketOptions(socketOptions)
          .build();
     LettuceClientConfiguration clientConfiguration = LettuceClientConfiguration.builder()
          .commandTimeout(Duration.ofMillis(redisReadTimeout))
          .readFrom(ReadFrom.MASTER)
          .clientOptions(clientOptions)
          .build();
     return clientConfiguration;
  }
}
Pooling configuration for single-node, master/standby, read/write splitting,
and Proxy Cluster instances
```

```
Enable the pooling component
<dependency>
<groupId>org.apache.commons</groupId>
<artifactId>commons-pool2</artifactId>
</dependency>
Code
```

import java.time.Duration;

import org.apache.commons.pool2.impl.GenericObjectPoolConfig; import org.springframework.beans.factory.annotation.Value; import org.springframework.context.annotation.Bean; import org.springframework.context.annotation.Configuration; import org.springframework.data.redis.connection.RedisConnectionFactory; import org.springframework.data.redis.connection.RedisStandaloneConfiguration; import org.springframework.data.redis.connection.lettuce.LettuceClientConfiguration; import org.springframework.data.redis.connection.lettuce.LettuceConnectionFactory; import org.springframework.data.redis.connection.lettuce.LettucePoolingClientConfiguration; import io.lettuce.core.ClientOptions; import io.lettuce.core.SocketOptions; \* Lettuce pooling configuration @Configuration public class RedisPoolConfiguration { @Value("\${redis.host}") private String redisHost; @Value("\${redis.port:6379}") private Integer redisPort = 6379; @Value("\${redis.database:0}") private Integer redisDatabase = 0; @Value("\${redis.password:}") private String redisPassword; @Value("\${redis.connect.timeout:2000}") private Integer redisConnectTimeout = 2000; @Value("\${redis.read.timeout:2000}") private Integer redisReadTimeout = 2000; @Value("\${redis.pool.minSize:50}") private Integer redisPoolMinSize = 50; @Value("\${redis.pool.maxSize:200}") private Integer redisPoolMaxSize = 200; @Value("\${redis.pool.maxWaitMillis:2000}") private Integer redisPoolMaxWaitMillis = 2000; @Value("\${redis.pool.softMinEvictableIdleTimeMillis:1800000}") private Integer redisPoolSoftMinEvictableIdleTimeMillis = 30 \* 60 \* 1000; @Value("\${redis.pool.timeBetweenEvictionRunsMillis:60000}") private Integer redisPoolBetweenEvictionRunsMillis = 60 \* 1000; \* TCP\_KEEPALIVE configuration parameters: A keepalive interval = TCP\_KEEPALIVE\_TIME = 30 Idle duration before keepalive = TCP\_KEEPALIVE\_TIME/3 = 10 \* \* keepalive xx times before disconnect = TCP\_KEEPALIVE\_COUNT = 3 \*/ private static final int TCP\_KEEPALIVE\_TIME = 30; \* TCP\_USER\_TIMEOUT Idle duration limit, to address Lettuce timeout. \* refer: https://github.com/lettuce-io/lettuce-core/issues/2082 \*/ private static final int TCP\_USER\_TIMEOUT = 30; @Bean public RedisConnectionFactory redisConnectionFactory(LettuceClientConfiguration clientConfiguration) {

RedisStandaloneConfiguration standaloneConfiguration = new RedisStandaloneConfiguration();

```
standaloneConfiguration.setHostName(redisHost);
     standaloneConfiguration.setPort(redisPort);
     standaloneConfiguration.setDatabase(redisDatabase);
     standaloneConfiguration.setPassword(redisPassword);
     LettuceConnectionFactory connectionFactory = new
LettuceConnectionFactory(standaloneConfiguration, clientConfiguration);
     connectionFactory.setDatabase(redisDatabase);
     //Disable sharing native connection before enabling pooling
     connectionFactory.setShareNativeConnection(false);
     return connectionFactory;
  }
  @Bean
  public LettuceClientConfiguration clientConfiguration() {
     SocketOptions socketOptions = SocketOptions.builder()
        .keepAlive(SocketOptions.KeepAliveOptions.builder()
          // A keepalive interval
          .idle(Duration.ofSeconds(TCP_KEEPALIVE_TIME))
          // Idle duration before keepalive
          .interval(Duration.ofSeconds(TCP_KEEPALIVE_TIME/3))
          // keepalive xx times before disconnect
          .count(3)
          // Whether to keep connections alive.
          .enable()
          .build())
        .tcpUserTimeout(SocketOptions.TcpUserTimeoutOptions.builder()
          // Addressing timeouts caused by RST on the server
          .tcpUserTimeout(Duration.ofSeconds(TCP_USER_TIMEOUT))
          .enable()
          .build())
        // TCP connection timeout setting
        .connectTimeout(Duration.ofMillis(redisConnectTimeout))
        .build();
     ClientOptions clientOptions = ClientOptions.builder()
          .autoReconnect(true)
          .pingBeforeActivateConnection(true)
          .cancelCommandsOnReconnectFailure(false)
          .disconnectedBehavior(ClientOptions.DisconnectedBehavior.ACCEPT_COMMANDS)
          .socketOptions(socketOptions)
          .build();
     LettucePoolingClientConfiguration clientConfiguration =
LettucePoolingClientConfiguration.builder()
          .poolConfig(poolConfig())
          .commandTimeout(Duration.ofMillis(redisReadTimeout))
          .clientOptions(clientOptions)
          .readFrom(ReadFrom.MASTER)
          .build();
     return poolingClientConfiguration;
  }
  private GenericObjectPoolConfig redisPoolConfig() {
     GenericObjectPoolConfig poolConfig = new GenericObjectPoolConfig();
     //Minimum idle connections in the pool
     poolConfig.setMinIdle(redisPoolMinSize);
     //Maximum idle connections in the pool
     poolConfig.setMaxIdle(redisPoolMaxSize);
     //Maximum total connections in the pool
     poolConfig.setMaxTotal(redisPoolMaxSize);
     //Wait when pool is exhausted? Set to true to wait. To validate setMaxWait, it has to be true.
     poolConfig.setBlockWhenExhausted(true);
     //Max allowed time to wait for connection after pool is exhausted. The default value -1 indicates
to wait indefinitely.
```

```
poolConfig.setMaxWait(Duration.ofMillis(redisPoolMaxWaitMillis));
```

<pre>//Set to true to enable connectivity test on creating connections. Default: false. poolConfig.setTestOnCreate(false);</pre>	
<pre>//Set to true to enable connectivity test on borrowing connections. Default: false. Set heavy-traffic services to reduce overhead. poolConfig.setTestOnBorrow(true):</pre>	to false for
//Set to true to enable connectivity test on returning connections. Default: false. Set to heavy-traffic services to reduce overhead.	o false for
poolConfig.setTestOnReturn(false); //Indicates whether to check for idle connections. If this is set to false, idle connection evicted.	s are not
poolConfig.setTestWhileIdle(true); //Idle duration after which a connection is evicted. If the actual duration is greater that value and the maximum number of idle connections is reached, idle connections are directle	an this y evicted.
<pre>poolConfig.setSoftMinEvictableIdleTime(Duration.ofMillis(redisPoolSoftMinEvictableIdleTim //Disable eviction policy MinEvictableIdleTimeMillis(). poolConfig.setMinEvictableIdleTime(Duration.ofMillis(-1)); //Interval for checking and evicting idle connections. Default: 60s. poolConfig.setTimeBetweenEvictionRuns(Duration.ofMillis(redisPoolBetweenEvictionR return poolConfig; }</pre>	eMillis)); unsMillis));
}	
Configuration for Redis Cluster instances     import iava time Duration:	
import java.util.ArrayList; import java.util.List;	
import org.springframework.beans.factory.annotation.Value; import org.springframework.context.annotation.Bean; import org.springframework.context annotation.Configuration;	
import org.springframework.data.redis.connection.RedisClusterConfiguration;	
import org.springframework.data.redis.connection.RedisConnectionFactory; import org.springframework.data.redis.connection.RedisNode;	
import org.springframework.data.redis.connection.lettuce.LettuceClientConfiguration; import org.springframework.data.redis.connection.lettuce.LettuceConnectionFactory;	
import io.lettuce.core.ClientOptions; import io.lettuce.core.SocketOptions;	
import io.lettuce.core.cluster.ClusterClientOptions; import io.lettuce.core.cluster.ClusterTopologyRefreshOptions;	
/** * Lettuce Cluster non-pooling configuration (use either this or the application properties co	nfiguration)
*/ @Configuration	gu.u.e.,
public class RedisConfiguration {	
@Value("\${redis.cluster.nodes}") private String redisClusterNodes;	
<pre>@Value("\${redis.cluster.maxDirects:3}") private Integer redisClusterMaxDirects;</pre>	
@Value("\${redis.password:}") private String redisPassword;	
@Value("\${redis.connect.timeout:2000}") private Integer redisConnectTimeout = 2000;	
@Value("\${redis.read.timeout:2000}") private Integer redisReadTimeout = 2000;	
@Value("\${redis.cluster.topology.refresh.period.millis:10000}") private Integer redisClusterTopologyRefreshPeriodMillis = 10000; /**	
<ul> <li>* TCP_KEEPALIVE configuration parameters:</li> <li>* A learnable internal TCP_KEEPALINE TIME _ 20</li> </ul>	
<ul> <li>A keepalive interval = ICP_KEEPALIVE_TIME = 30</li> <li>* Idle duration before keepalive = TCP_KEEPALIVE_TIME/3 = 10</li> </ul>	
* keepalive xx times before disconnect = TCP_KEEPALIVE_COUNT = 3	

```
private static final int TCP_KEEPALIVE_TIME = 30;
   * TCP_USER_TIMEOUT Idle duration limit, to address Lettuce timeout.
   * refer: https://github.com/lettuce-io/lettuce-core/issues/2082
   */
  private static final int TCP_USER_TIMEOUT = 30;
  @Bean
  public RedisConnectionFactory redisConnectionFactory(LettuceClientConfiguration
clientConfiguration) {
     RedisClusterConfiguration clusterConfiguration = new RedisClusterConfiguration();
     List<RedisNode> clusterNodes = new ArrayList<>();
     for (String clusterNodeStr : redisClusterNodes.split(",")) {
        String[] nodeInfo = clusterNodeStr.split(":");
        clusterNodes.add(new RedisNode(nodeInfo[0], Integer.valueOf(nodeInfo[1])));
     }
     clusterConfiguration.setClusterNodes(clusterNodes);
     clusterConfiguration.setPassword(redisPassword);
     clusterConfiguration.setMaxRedirects(redisClusterMaxDirects);
     LettuceConnectionFactory connectionFactory = new
LettuceConnectionFactory(clusterConfiguration, clientConfiguration);
     return connectionFactory;
  }
  @Bean
  public LettuceClientConfiguration clientConfiguration() {
     SocketOptions socketOptions = SocketOptions.builder()
        .keepAlive(SocketOptions.KeepAliveOptions.builder()
          // A keepalive interval
          .idle(Duration.ofSeconds(TCP_KEEPALIVE_TIME))
          // Idle duration before keepalive
          .interval(Duration.ofSeconds(TCP_KEEPALIVE_TIME/3))
          // keepalive xx times before disconnect
          .count(3)
          // Whether to keep connections alive.
          .enable()
          .build())
        .tcpUserTimeout(SocketOptions.TcpUserTimeoutOptions.builder()
          // Addressing timeouts caused by RST on the server
          .tcpUserTimeout(Duration.ofSeconds(TCP_USER_TIMEOUT))
          .enable()
          .build())
        // TCP connection timeout setting
        .connectTimeout(Duration.ofMillis(redisConnectTimeout))
        .build();
     ClusterTopologyRefreshOptions topologyRefreshOptions =
ClusterTopologyRefreshOptions.builder()
          .enableAllAdaptiveRefreshTriggers()
          .enablePeriodicRefresh(Duration.ofMillis(redisClusterTopologyRefreshPeriodMillis))
          .build();
     ClusterClientOptions clientOptions = ClusterClientOptions.builder()
          .autoReconnect(true)
          .pingBeforeActivateConnection(true)
          .cancelCommandsOnReconnectFailure(false)
          .disconnectedBehavior(ClientOptions.DisconnectedBehavior.ACCEPT_COMMANDS)
          .socketOptions(socketOptions)
          .topologyRefreshOptions(topologyRefreshOptions)
          .build();
```

LettuceClientConfiguration clientConfiguration = LettuceClientConfiguration.builder()

Issue 01 (2025-04-27) Copyright © Huawei Cloud Computing Technologies Co., Ltd.

```
.commandTimeout(Duration.ofMillis(redisReadTimeout))
          .readFrom(ReadFrom.MASTER)
          .clientOptions(clientOptions)
          .build();
     return clientConfiguration;
  }
}
Pooling configuration for Redis Cluster instances
Enable the pooling component
<dependency>
  <groupId>org.apache.commons</groupId>
  <artifactId>commons-pool2</artifactId>
   <version>2.11.1</version>
</dependency>
Code
import java.time.Duration;
import java.util.ArrayList;
import java.util.List;
import org.apache.commons.pool2.impl.GenericObjectPoolConfig;
import org.springframework.beans.factory.annotation.Value;
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.Configuration;
import org.springframework.data.redis.connection.RedisClusterConfiguration;
import org.springframework.data.redis.connection.RedisConnectionFactory;
import org.springframework.data.redis.connection.RedisNode;
import org.springframework.data.redis.connection.lettuce.LettuceClientConfiguration;
import org.springframework.data.redis.connection.lettuce.LettuceConnectionFactory;
import org.springframework.data.redis.connection.lettuce.LettucePoolingClientConfiguration;
import io.lettuce.core.ClientOptions;
import io.lettuce.core.SocketOptions;
import io.lettuce.core.cluster.ClusterClientOptions;
import io.lettuce.core.cluster.ClusterTopologyRefreshOptions;
 Lettuce pooling configuration
@Configuration
public class RedisPoolConfiguration {
  @Value("${redis.cluster.nodes}")
  private String redisClusterNodes;
  @Value("${redis.cluster.maxDirects:3}")
  private Integer redisClusterMaxDirects;
  @Value("${redis.password:}")
  private String redisPassword;
  @Value("${redis.connect.timeout:2000}")
  private Integer redisConnectTimeout = 2000;
  @Value("${redis.read.timeout:2000}")
  private Integer redisReadTimeout = 2000;
  @Value("${redis.cluster.topology.refresh.period.millis:10000}")
  private Integer redisClusterTopologyRefreshPeriodMillis = 10000;
  @Value("${redis.pool.minSize:50}")
  private Integer redisPoolMinSize = 50;
  @Value("${redis.pool.maxSize:200}")
  private Integer redisPoolMaxSize = 200;
  @Value("${redis.pool.maxWaitMillis:2000}")
  private Integer redisPoolMaxWaitMillis = 2000;
```

```
@Value("${redis.pool.softMinEvictableIdleTimeMillis:1800000}")
  private Integer redisPoolSoftMinEvictableIdleTimeMillis = 30 * 60 * 1000;
  @Value("${redis.pool.timeBetweenEvictionRunsMillis:60000}")
  private Integer redisPoolBetweenEvictionRunsMillis = 60 * 1000;
   * TCP_KEEPALIVE configuration parameters:
   *
    A keepalive interval = TCP_KEEPALIVE_TIME = 30
     Idle duration before keepalive = TCP_KEEPALIVE_TIME/3 = 10
     keepalive xx times before disconnect = TCP_KEEPALIVE_COUNT = 3
  private static final int TCP_KEEPALIVE_TIME = 30;
   * TCP_USER_TIMEOUT Idle duration limit, to address Lettuce timeout.
   * refer: https://github.com/lettuce-io/lettuce-core/issues/2082
   */
  private static final int TCP_USER_TIMEOUT = 30;
  @Bean
  public RedisConnectionFactory redisConnectionFactory(LettuceClientConfiguration
clientConfiguration) {
     RedisClusterConfiguration clusterConfiguration = new RedisClusterConfiguration();
     List<RedisNode> clusterNodes = new ArrayList<>();
     for (String clusterNodeStr : redisClusterNodes.split(",")) {
        String[] nodeInfo = clusterNodeStr.split(":");
        clusterNodes.add(new RedisNode(nodeInfo[0], Integer.valueOf(nodeInfo[1])));
     }
     clusterConfiguration.setClusterNodes(clusterNodes);
     clusterConfiguration.setPassword(redisPassword);
     clusterConfiguration.setMaxRedirects(redisClusterMaxDirects);
     LettuceConnectionFactory connectionFactory = new
LettuceConnectionFactory(clusterConfiguration, clientConfiguration);
     //Disable native connection sharing before validating connection pool
     connectionFactory.setShareNativeConnection(false);
     return connectionFactory;
  }
  @Bean
  public LettuceClientConfiguration clientConfiguration() {
     SocketOptions socketOptions = SocketOptions.builder()
        .keepAlive(SocketOptions.KeepAliveOptions.builder()
          // A keepalive interval
          .idle(Duration.ofSeconds(TCP_KEEPALIVE_TIME))
          // Idle duration before keepalive
          .interval(Duration.ofSeconds(TCP KEEPALIVE TIME/3))
          // keepalive xx times before disconnect
          .count(3)
          // Whether to keep connections alive.
          .enable()
          .build())
        .tcpUserTimeout(SocketOptions.TcpUserTimeoutOptions.builder()
          // Addressing timeouts caused by RST on the server
          .tcpUserTimeout(Duration.ofSeconds(TCP_USER_TIMEOUT))
          .enable()
          .build())
        // TCP connection timeout setting
        .connectTimeout(Duration.ofMillis(redisConnectTimeout))
        .build();
     ClusterTopologyRefreshOptions topologyRefreshOptions =
ClusterTopologyRefreshOptions.builder()
          .enableAllAdaptiveRefreshTriggers()
```

.build(),	
ClusterClientOptions clientOptions = ClusterClientOptions.builder() .autoReconnect(true) .pingBeforeActivateConnection(true) .cancelCommandsOnReconnectFailure(false) .disconnectedBehavior(ClientOptions.DisconnectedBehavior.ACCEPT_COMMANDS) .socketOptions(socketOptions) .topologyRefreshOptions(topologyRefreshOptions) .build();	
LettucePoolingClientConfiguration clientConfiguration = LettucePoolingClientConfiguration.builder() .poolConfig(poolConfig()) .commandTimeout(Duration.ofMillis(redisReadTimeout)) .clientOptions(clientOptions) .readFrom(ReadFrom.MASTER) .build(); return clientConfiguration; }	
<pre>private GenericObjectPoolConfig poolConfig() {     GenericObjectPoolConfig poolConfig = new GenericObjectPoolConfig();     //Minimum connections in the pool     poolConfig.setMinIdle(redisPoolMinSize);     //Maximum idle connections in the pool     poolConfig.setMaxIdle(redisPoolMaxSize);     //Maximum total connections in the pool     poolConfig.setMaxTotal(redisPoolMaxSize);     //Maximum total connections in the pool     poolConfig.setMaxTotal(redisPoolMaxSize);     //Maxit when pool is exhausted? Set to true to wait. To validate setMaxWait, it has to be tru     poolConfig.setBlockWhenExhausted(true);     //Max allowed time to wait for connection after pool is exhausted. The default value -1 ind     to wait indefinitely.     poolConfig.setTaxWait(Duration.ofMillis(redisPoolMaxWaitMillis));     //Set to true to enable connectivity test on creating connections. Default: false.     poolConfig.setTestOnCreate(false);     //Set to true to enable connectivity test on borrowing connections. Default: false. Set to fals     heavy-traffic services to reduce overhead.     poolConfig.setTestOnBorrow(true);     //Set to true to enable connectivity test on returning connections. Default: false. Set to false heavy-traffic services to reduce overhead.     poolConfig.setTestOnBorrow(true);     //Set to true to enable connectivity test on returning connections. Default: false. Set to false heavy-traffic services to reduce overhead.     poolConfig.setTestOnReturn(false);     //Indicates whether to check for idle connections. If this is set to false, idle connections are evicted.     poolConfig.setTestWhileIdle(true);     //Disable connection closure when the minimum idle time is reached.     poolConfig.setTestWhileIdle(true);     //Disable connection closure when the minimum idle time is reached.     poolConfig.setTestWhileIdle(Time(Duration.ofMillis(-1));     //Idle duration before a connection being evicted. If the actual duration is greater than this and the maximum number of idle connections is reached, idle connectio</pre>	e. icates se for e for not value s));
poolConfig.setSoftMinEvictableIdleTime(Duration.ofMillis(redisPoolSoftMinEvictableIdleTimeMilli //Interval for checking and evicting idle connections. Default: 60s. poolConfig.setTimeBetweenEvictionRuns(Duration.ofMillis(redisPoolBetweenEvictionRunsM return poolConfig;	s)); illis));
}	

# (Optional) Configuring SSL Connections

If SSL is enabled for an instance, to access it using SSL connections, use the following content to replace the **LettuceClientConfiguration** construction method **clientConfiguration()** in **Bean Configuration**. For details about whether your

#### DCS Redis instances support SSL, see **Transmitting DCS Redis Data with Encryption Using SSL**.

Single-node, master/standby, read/write splitting, and Proxy Cluster @Bean public LettuceClientConfiguration clientConfiguration() { SocketOptions socketOptions = SocketOptions.builder() .keepAlive(SocketOptions.KeepAliveOptions.builder() // A keepalive interval .idle(Duration.ofSeconds(TCP\_KEEPALIVE\_TIME)) // Idle duration before keepalive .interval(Duration.ofSeconds(TCP\_KEEPALIVE\_TIME/3)) // keepalive xx times before disconnect .count(3) // Whether to keep connections alive. .enable() .build()) .tcpUserTimeout(SocketOptions.TcpUserTimeoutOptions.builder() // Addressing timeouts caused by RST on the server .tcpUserTimeout(Duration.ofSeconds(TCP\_USER\_TIMEOUT)) .enable() .build()) // TCP connection timeout setting .connectTimeout(Duration.ofMillis(redisConnectTimeout)) .build(); SslOptions sslOptions = SslOptions.builder() .trustManager(new File(certificationPath)) .build(): ClientOptions clientOptions = ClientOptions.builder() .sslOptions(sslOptions) .autoReconnect(true) .pingBeforeActivateConnection(true) .cancelCommandsOnReconnectFailure(false) .disconnectedBehavior(ClientOptions.DisconnectedBehavior.ACCEPT\_COMMANDS) .socketOptions(socketOptions) .build(); LettuceClientConfiguration clientConfiguration = LettuceClientConfiguration.builder() .commandTimeout(Duration.ofMillis(redisReadTimeout)) .readFrom(ReadFrom.MASTER) .clientOptions(clientOptions) .useSsl() .build(); return clientConfiguration; } **Redis Cluster** @Bean public LettuceClientConfiguration clientConfiguration() { SocketOptions socketOptions = SocketOptions.builder() .keepAlive(SocketOptions.KeepAliveOptions.builder() // A keepalive interval .idle(Duration.ofSeconds(TCP\_KEEPALIVE\_TIME)) // Idle duration before keepalive .interval(Duration.ofSeconds(TCP\_KEEPALIVE\_TIME/3)) // keepalive xx times before disconnect .count(3) // Whether to keep connections alive. .enable() .build()) .tcpUserTimeout(SocketOptions.TcpUserTimeoutOptions.builder() // Addressing timeouts caused by RST on the server .tcpUserTimeout(Duration.ofSeconds(TCP\_USER\_TIMEOUT)) .enable() .build()) // TCP connection timeout setting .connectTimeout(Duration.ofMillis(redisConnectTimeout))

.build();

```
SslOptions sslOptions = SslOptions.builder()
     .trustManager(new File(certificationPath))
     .build();
  ClusterTopologyRefreshOptions topologyRefreshOptions = ClusterTopologyRefreshOptions.builder()
     .enableAllAdaptiveRefreshTriggers()
     .enablePeriodicRefresh(Duration.ofMillis(redisClusterTopologyRefreshPeriodMillis))
     .build();
  ClusterClientOptions clientOptions = ClusterClientOptions.builder()
     .sslOptions(sslOptions)
     .autoReconnect(true)
     .pingBeforeActivateConnection(true)
     .cancelCommandsOnReconnectFailure(false)
     .disconnectedBehavior(ClientOptions.DisconnectedBehavior.ACCEPT_COMMANDS)
     .socketOptions(socketOptions)
     .topologyRefreshOptions(topologyRefreshOptions)
     .build();
  LettuceClientConfiguration clientConfiguration = LettuceClientConfiguration.builder()
     .commandTimeout(Duration.ofMillis(redisReadTimeout))
     .readFrom(ReadFrom.MASTER)
     .clientOptions(clientOptions)
     .useSsl()
     .build();
  return clientConfiguration;
}
```

## **Parameter Description**

Parameter	Туре	Default Value	Description
configuration	RedisConfi guration	-	<ul><li>Redis connection configuration.</li><li>Two subsclasses:</li><li>RedisStandaloneConfiguration</li><li>RedisClusterConfiguration</li></ul>
clientConfigur ation	LettuceCli entConfig uration	-	Client configuration parameter. Common subclass: LettucePoolingClientConfiguration
shareNativeCo nnection	boolean	true	Indicates whether to share native connections. Set to <b>true</b> to share. Set to <b>false</b> to enable connection pooling.

Table 4-7 LettuceConnectionFactory pa	arameters
---------------------------------------	-----------

Parameter	Default Value	Description
hostName	localhost	IP address/domain name for connecting to a DCS Redis instance
port	6379	Port number
database	0	Database subscript
password	-	Password

 Table 4-8 RedisStandaloneConfiguration parameters

#### Table 4-9 RedisClusterConfiguration parameters

Parameter	Description
clusterNodes	Cluster node connection information, including the node IP address and port number
maxRedirects	Maximum redirecting times. Recommended value: <b>3</b> .
password	Password

 Table 4-10 LettuceClientConfiguration parameters

Parameter	Туре	Default Value	Description
timeout	Duration	60s	Command timeout: Recommended: <b>2s</b> .
clientOptions	ClientOpti ons	-	Configuration options.
readFrom	readFrom	MASTER	Read mode. Recommended: <b>MASTER</b> . Other values may cause access failures in failover scenarios.

Parameter	Туре	Default Value	Description
timeout	Duration	60s	Command timeout: Recommended: <b>2s</b> .
clientOptions	ClientOpti ons	-	Configuration options.

Parameter	Туре	Default Value	Description
poolConfig	GenericOb jectPoolCo nfig	-	Connection pool configuration.
readFrom	readFrom	MASTER	Read mode. Recommended: <b>MASTER</b> . Other values may cause access failures in failover scenarios.

 Table 4-12 ClientOptions parameters

Parameter	Туре	Default Value	Description
autoReconnect	boolean	true	Indicates whether to automatically reconnect after disconnection. Recommended: <b>true</b> .
pingBeforeActi vateConnectio n	boolean	true	Indicates whether to test connectivity on established connections. Recommended: <b>true</b> .
cancelComma ndsOnReconne ctFailure	boolean	true	Indicates whether to cancel commands after a failed reconnection attempt. Recommended: <b>false</b> .
disconnectedB ehavior	Disconnec tedBehavi or	Disconnecte dBehavior.DE FAULT	<ul> <li>Indicates what to do when a connection drops. Recommended:</li> <li>ACCEPT_COMMANDS.</li> <li>DEFAULT: When autoReconnect is set true, commands are allowed to wait in queue. When autoReconnect is set to false, commands are not allowed to wait in queue.</li> <li>ACCEPT_COMMANDS: Allow commands to wait in queue.</li> <li>REJECT_COMMANDS: Do not allow commands to wait in queue.</li> </ul>
socketOptions	SocketOpt ions	-	Socket configuration.

Parameter	Default Value	Description
connectTimeo ut	10s	Connection timeout. Recommended: <b>2s</b> .

## Table 4-14 GenericObjectPoolConfig parameters

Parameter	Default Value	Description
minIdle	-	Minimum connections in the pool.
maxIdle	-	Maximum idle connections in the connection pool.
maxTotal	-	Maximum total connections in the connection pool.
blockWhenExha usted	true	Indicates whether to wait after the connection pool is exhausted. <b>true</b> : Wait. <b>false</b> : Do not wait. To validate <b>maxWaitMillis</b> , this parameter must be set to <b>true</b> .
maxWaitMillis	-1	Maximum amount of time a connection allocation should block before throwing an exception when the pool is exhausted. The default value <b>-1</b> indicates to wait indefinitely.
testOnCreate	false	Set to true to enable connectivity test on creating connections. Default: <b>false</b> .
testOnBorrow	false	Set to true to enable connectivity test on borrowing connections. Default: <b>false</b> . Set to false for heavy-traffic services to reduce overhead.
testOnReturn	false	Set to <b>true</b> to enable connectivity test on returning connections. Default: <b>false</b> . Set to <b>false</b> for heavy-traffic services to reduce overhead.
testWhileIdle	false	Indicates whether to check for idle connections. If this parameter is set to <b>false</b> , idle connections are not evicted. Recommended value: <b>true</b> .

Parameter	Default Value	Description
softMinEvictabl eIdleTimeMillis	-1	Duration (in milliseconds) after which idle connections are evicted. If the idle duration is greater than this value and the maximum number of idle connections is reached, idle connections are directly evicted. Recommended value: <b>1800000</b> .
minEvictableIdle TimeMillis	1800000	An eviction policy, set to <b>-1</b> (suggested) to disable it. Use <b>softminEvictableIdleTimeMil-lis</b> instead.
timeBetweenEvi ctionRunsMillis	-1	Eviction interval, in milliseconds. Recommended value: <b>60000</b>

## **Suggestion for Configuring DCS Instances**

• Pooling connection

Different from Jedis's BIO, the bottom layer of Lettuce communicates with Redis Server based on Netty's NIO. Combining persistent connections and queues, Lettuce sends and receives multiple requests and responses spontaneously with sequential sending and receiving features of TCP. A single connection supports 3000 to 5000 QPS, but you are not advised to allow more than 3000 QPS in production systems. Pooling is not supported by Lettuce, and is disabled by default in Spring Boot. To enable pooling, validate the commons-pool2 dependency and disable native connection sharing.

By default, each Lettuce connection needs two thread pools, I/O thread pool and computation thread pool, to support I/O event reading and asynchronous event processing. If you configure connection pooling, each connection creates two thread pools, consuming high memory resources. Lettuce is strong at processing single connections based on its bottom-layer implementation, so you are not advised to use Lettuce with pooling.

Topology refresh

When connecting to a Redis Cluster instance, Lettuce randomly sends **cluster nodes** to the node list during initialization to obtain the distribution of cluster slots. Cluster topology structure changes when the cluster capacity is increased or decreased or a master/standby switchover occurs. Lettuce does not detect such changes by default. You can enable detection with the following configurations:

application.properties configuration
 #Enable adaptive topology refresh.
 spring.redis.lettuce.cluster.refresh.adaptive=true
 #Enable topology refresh every 10 seconds.
 spring.redis.lettuce.cluster.refresh.period=10S
 API configuration
 ClusterTopologyRefreshOptions topologyRefreshOptions =
 ClusterTopologyRefreshOptions.builder()
 .enableAllAdaptiveRefreshTriggers()
 .enablePeriodicRefresh(Duration.ofMillis(redisClusterTopologyRefreshPeriodMillis))
 .build();

ClusterClientOptions clientOptions = ClusterClientOptions.builder()

```
...
.topologyRefreshOptions(topologyRefreshOptions)
.build();
```

Blast radius

The bottom layer of Lettuce uses a combination of single persistent connection and request queue. Once network jitter or intermittent disconnection occurs or connection times out, all requests are affected. Especially when connection times out, an attempt is made to resend TCP pockets until timeout and connection drops. Requests do not recover until connections are reestablished. Requests accumulate during resending attempts. If upper-layer services time out in batches, or the resending timeout is too long in some OSs' kernels, the service system remains unavailable for a long time. **Therefore, you are advised to use Jedis instead of Lettuce.** 

# 4.3.4 Connecting to Redis on Redisson (Java)

This section describes how to access a Redis instance on Redisson. For more information about how to use other Redis clients, visit **the Redis official website**.

For Spring Boot projects, Spring Data Redis is already integrated with Jedis and Lettuce, but does not support Redisson. Redisson provides the redisson-spring-boot-starter component (https://mvnrepository.com/artifact/org.redisson/redisson) that can be used with Spring Boot.

Spring Boot 1.x is integrated with Jedis, and Spring Boot 2.x is integrated with Lettuce.

#### Notes and Constraints

- If a password was set during DCS Redis instance creation, configure the password for connecting to Redis using Redisson. Do not hard code the plaintext password.
- To connect to a single-node, read/write splitting, or Proxy Cluster instance, use the useSingleServer method of the SingleServerConfig object of Redisson. To connect to a master/standby instance, use the useMasterSlaveServers method of the MasterSlaveServersConfig object of Redisson. To connect to a Redis Cluster instance, use the useClusterServers method of the ClusterServersConfig object.
- Springboot 2.3.12.RELEASE or later is required. Redisson **3.37.0** or later is required.

## Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- View the IP address/domain name and port of the DCS Redis instance to be accessed. For details, see Viewing and Modifying Basic Settings of a DCS Instance.

## Pom Configuration

<!-- spring-data-redis -->

```
<groupId>org.springframework.boot</groupId>
  <artifactId>spring-boot-starter-data-redis</artifactId>
  <exclusions>
     <!--Lettuce is integrated in Spring Boot 2.x by default. This dependency needs to be deleted. -->
     <exclusion>
        <artifactId>lettuce-core</artifactId>
        <groupId>io.lettuce</groupId>
     </exclusion>
  </exclusions>
</dependency>
<!--Redisson's adaptation package for Spring Boot-->
<dependency>
  <groupId>org.redisson</groupId>
  <artifactId>redisson-spring-boot-starter</artifactId>
  <version>${redisson.version}</version>
</dependency>
```

#### **Bean Configuration**

Spring Boot does not provide Redisson adaptation, and the **application.properties** configuration file does not have the corresponding configuration item. Therefore, you can only use Bean configuration.

```
Single-node, read/write splitting, and Proxy Cluster
import org.redisson.Redisson;
import org.redisson.api.RedissonClient;
import org.redisson.codec.JsonJacksonCodec;
import org.redisson.config.Config;
import org.redisson.config.SingleServerConfig;
import org.springframework.beans.factory.annotation.Value;
import org.springframework.context.annotation.Bean;
import org.springframework.context.annotation.Configuration;
@Configuration
public class SingleConfig {
  @Value("${redis.address:}")
  private String redisAddress;
  @Value("${redis.password:}")
  private String redisPassword;
  @Value("${redis.database:0}")
  private Integer redisDatabase = 0;
  @Value("${redis.connect.timeout:3000}")
  private Integer redisConnectTimeout = 3000;
  @Value("${redis.connection.idle.timeout:10000}")
  private Integer redisConnectionIdleTimeout = 10000;
  @Value("${redis.connection.ping.interval:1000}")
  private Integer redisConnectionPingInterval = 1000;
  @Value("${redis.timeout:2000}")
  private Integer timeout = 2000;
  @Value("${redis.connection.pool.min.size:50}")
  private Integer redisConnectionPoolMinSize;
  @Value("${redis.connection.pool.max.size:200}")
  private Integer redisConnectionPoolMaxSize;
  @Value("${redis.retry.attempts:3}")
  private Integer redisRetryAttempts = 3;
  @Value("${redis.retry.interval:200}")
  private Integer redisRetryInterval = 200;
```

#### @Bean

public RedissonClient redissonClient(){
 Config redissonConfig = new Config();

SingleServerConfig serverConfig = redissonConfig.useSingleServer(); serverConfig.setAddress(redisAddress); serverConfig.setConnectionMinimumIdleSize(redisConnectionPoolMinSize); serverConfig.setConnectionPoolSize(redisConnectionPoolMaxSize);

serverConfig.setDatabase(redisDatabase); serverConfig.setPassword(redisPassword); serverConfig.setConnectTimeout(redisConnectTimeout); serverConfig.setIdleConnectionTimeout(redisConnectionIdleTimeout); serverConfig.setPingConnectionInterval(redisConnectionPingInterval); serverConfig.setTimeout(timeout); serverConfig.setRetryAttempts(redisRetryAttempts); serverConfig.setRetryInterval(redisRetryInterval);

redissonConfig.setCodec(new JsonJacksonCodec()); return Redisson.create(redissonConfig);

```
}
}
```

#### Master/Standby

import org.redisson.Redisson; import org.redisson.api.RedissonClient; import org.redisson.codec.JsonJacksonCodec; import org.redisson.config.Config; import org.redisson.config.MasterSlaveServersConfig; import org.redisson.config.ReadMode; import org.redisson.config.SubscriptionMode; import org.springframework.beans.factory.annotation.Value; import org.springframework.context.annotation.Bean; import org.springframework.context.annotation.Configuration; import java.util.HashSet; @Configuration public class MasterStandbyConfig { @Value("\${redis.master.address}") private String redisMasterAddress; @Value("\${redis.slave.address}") private String redisSlaveAddress; @Value("\${redis.database:0}") private Integer redisDatabase = 0; @Value("\${redis.password:}") private String redisPassword; @Value("\${redis.connect.timeout:3000}") private Integer redisConnectTimeout = 3000; @Value("\${redis.connection.idle.timeout:10000}") private Integer redisConnectionIdleTimeout = 10000; @Value("\${redis.connection.ping.interval:1000}") private Integer redisConnectionPingInterval = 1000; @Value("\${redis.timeout:2000}") private Integer timeout = 2000; @Value("\${redis.master.connection.pool.min.size:50}") private Integer redisMasterConnectionPoolMinSize = 50; @Value("\${redis.master.connection.pool.max.size:200}") private Integer redisMasterConnectionPoolMaxSize = 200;

@Value("\${redis.retry.attempts:3}")
private Integer redisRetryAttempts = 3;

@Value("\${redis.retry.interval:200}")
private Integer redisRetryInterval = 200;

@Bean

public RedissonClient redissonClient() {
 Config redissonConfig = new Config();

MasterSlaveServersConfig serverConfig = redissonConfig.useMasterSlaveServers(); serverConfig.setMasterAddress(redisMasterAddress); HashSet<String> slaveSet = new HashSet<>(); slaveSet.add(redisSlaveAddress); serverConfig.setSlaveAddresses(slaveSet);

serverConfig.setDatabase(redisDatabase);
serverConfig.setPassword(redisPassword);

serverConfig.setMasterConnectionMinimumIdleSize(redisMasterConnectionPoolMinSize); serverConfig.setMasterConnectionPoolSize(redisMasterConnectionPoolMaxSize);

serverConfig.setReadMode(ReadMode.MASTER); serverConfig.setSubscriptionMode(SubscriptionMode.MASTER);

serverConfig.setConnectTimeout(redisConnectTimeout); serverConfig.setIdleConnectionTimeout(redisConnectionIdleTimeout); serverConfig.setPingConnectionInterval(redisConnectionPingInterval); serverConfig.setTimeout(timeout); serverConfig.setRetryAttempts(redisRetryAttempts); serverConfig.setRetryInterval(redisRetryInterval);

redissonConfig.setCodec(new JsonJacksonCodec()); return Redisson.create(redissonConfig);

Redis Cluster

} }

import org.redisson.Redisson; import org.redisson.api.RedissonClient; import org.redisson.codec.JsonJacksonCodec; import org.redisson.config.ClusterServersConfig; import org.redisson.config.Config; import org.redisson.config.ReadMode; import org.redisson.config.SubscriptionMode; import org.springframework.beans.factory.annotation.Value; import org.springframework.context.annotation.Bean; import org.springframework.context.annotation.Configuration;

import java.util.List;

@Configuration
public class ClusterConfig {

@Value("\${redis.cluster.address}")
private List<String> redisClusterAddress;

@Value("\${redis.cluster.scan.interval:5000}")
private Integer redisClusterScanInterval = 5000;

@Value("\${redis.password:}")
private String redisPassword;

@Value("\${redis.connect.timeout:3000}")
private Integer redisConnectTimeout = 3000;

@Value("\${redis.connection.idle.timeout:10000}")
private Integer redisConnectionIdleTimeout = 10000;

@Value("\${redis.connection.ping.interval:1000}")

private Integer redisConnectionPingInterval = 1000;

@Value("\${redis.timeout:2000}")
private Integer timeout = 2000;

@Value("\${redis.retry.attempts:3}")
private Integer redisRetryAttempts = 3;

@Value("\${redis.retry.interval:200}")
private Integer redisRetryInterval = 200;

@Value("\${redis.master.connection.pool.min.size:50}")
private Integer redisMasterConnectionPoolMinSize = 50;

@Value("\${redis.master.connection.pool.max.size:200}")
private Integer redisMasterConnectionPoolMaxSize = 200;

@Bean

public RedissonClient redissonClient() {
 Config redissonConfig = new Config();

ClusterServersConfig serverConfig = redissonConfig.useClusterServers(); serverConfig.setNodeAddresses(redisClusterAddress); serverConfig.setScanInterval(redisClusterScanInterval);

serverConfig.setPassword(redisPassword);

serverConfig.setMasterConnectionMinimumIdleSize(redisMasterConnectionPoolMinSize); serverConfig.setMasterConnectionPoolSize(redisMasterConnectionPoolMaxSize);

serverConfig.setReadMode(ReadMode.MASTER);
serverConfig.setSubscriptionMode(SubscriptionMode.MASTER);

serverConfig.setConnectTimeout(redisConnectTimeout); serverConfig.setIdleConnectionTimeout(redisConnectionIdleTimeout); serverConfig.setPingConnectionInterval(redisConnectionPingInterval); serverConfig.setTimeout(timeout); serverConfig.setRetryAttempts(redisRetryAttempts); serverConfig.setRetryInterval(redisRetryInterval);

redissonConfig.setCodec(new JsonJacksonCodec()); return Redisson.create(redissonConfig);

}

## (Optional) Configuring SSL Connections

If SSL is enabled for an instance, to access it using SSL connections, add the **configRedissonSSL(serverConfig)** logic to the **RedissonClient** construction method **clientConfiguration()** in **Bean Configuration** and change the Redis addresses from **redis:**//*ip*.*port* to **rediss:**//*ip*.*port*. For details about whether your DCS Redis instances support SSL, see **Transmitting DCS Redis Data with Encryption Using SSL**.

```
private void configRedissonSSL(BaseConfig serverConfig) {
  TrustManagerFactory trustManagerFactory = null;
  try {
    //Load the CA certificate in the user-defined path.
    CertificateFactory cf = CertificateFactory.getInstance("X.509");
    Certificate ca;
    try (InputStream is = new FileInputStream(certificationPath)) {
      ca = cf.generateCertificate(is);
    }
    //Create keystore.
    String keyStoreType = KeyStore.getDefaultType();
    KeyStore keyStore = KeyStore.getInstance(keyStoreType);
```

```
keyStore.load(null, null);
keyStore.setCertificateEntry("ca", ca);
//Create TrustManager.
trustManagerFactory = TrustManagerFactory.getInstance(TrustManagerFactory.getDefaultAlgorithm());
trustManagerFactory.init(keyStore);
} catch (CertificateException | IOException | KeyStoreException | NoSuchAlgorithmException e) {
e.printStackTrace();
return;
}
serverConfig.setSslTrustManagerFactory(trustManagerFactory);
```

## **Parameter Description**

}

Parameter	Default Value	Description
codec	org.redisson.cod ec.JsonJacksonC odec	Encoding format, including JSON, Avro, Smile, CBOR, and MsgPack.
threads	Number of CPU cores x 2	Thread pool used for executing RTopic Listener, RRemoteService, and RExecutorService.
executor	null	The function is the same as <b>threads</b> . If this parameter is not set, a thread pool is initialized based on <b>threads</b> .
nettyThreads	Number of CPU cores x 2	Thread pool used by the TCP channel that connects to the redis-server. All channels share this connection pool and are mapped to Netty's <b>Bootstrap.group()</b> .
eventLoopGroup	null	The function is the same as <b>nettyThreads</b> . If this parameter is not set, an EventLoopGroup is initialized based on the <b>nettyThreads</b> parameter for the bottom- layer TCP channel to use.
transportMode	TransportMode. NIO	Transmission mode. The options are <b>NIO</b> , <b>EPOLL</b> (additional package required), and <b>KQUEUE</b> (additional package required).
lockWatchdogTi meout	30000	Timeout interval (in milliseconds) of the lock-monitoring watchdog. In the distributed lock scenario, if the <b>leaseTimeout</b> parameter is not specified, the default value of this parameter is used.
keepPubSubOrd er	true	Indicates whether to receive messages in the publish sequence. If messages can be processed concurrently, you are advised to set this parameter to false.

Table 4-15 Config parameters

Table 4-16 SingleServerConfig	parameters	(single-node,	read/write	splitting,	or
Proxy Cluster)					

Parameter	Default Value	Description
address	-	Node connection information, in redis:// <i>ip</i> : <i>port</i> format.
database	0	ID of the database to be used.
connectionMini mumIdleSize	32	Minimum number of connections to the master node of each shard.
connectionPoolS ize	64	Maximum number of connections to the master node of each shard.
subscriptionCon- nectionMinimu mIdleSize	1	Minimum number of connections to the target node for pub/sub.
subscriptionCon- nectionPoolSize	50	Maximum number of connections to the target node for pub/sub.
subcriptionPerCo nnection	5	Maximum number of subscriptions on each subscription connection.
connectionTime out	10000	Connection timeout interval, in milliseconds.
idleConnectionTi meout	10000	Maximum time (in milliseconds) for reclaiming idle connections.
pingConnectionI nterval	30000	Heartbeat for detecting available connections, in milliseconds. <b>Recommended: 3000 ms</b> .
timeout	3000	Timeout interval for waiting for a response, in milliseconds.
retryAttempts	3	Maximum number of retries upon send failures.
retryInterval	1500	Retry interval, in milliseconds. <b>Recommended: 200 ms</b> .
clientName	null	Client name.

Table 4-17 MasterSlaveServersConfig	parameters	(master/standby)
-------------------------------------	------------	------------------

Parameter	Default Value	Description
masterAddress	-	Master node connection information, in redis:// <i>ip.port</i> format.

Parameter	Default Value	Description
slaveAddresses	-	Standby node connection information, in Set <redis: <i="">ip:port&gt; format.</redis:>
readMode	SLAVE	Read mode. By default, read traffic is distributed to replica nodes. The value can be <b>MASTER</b> (recommended), <b>SLAVE</b> , or <b>MASTER_SLAVE</b> . Other values may cause access failures in failover scenarios.
loadBalancer	RoundRobinLoad Balancer	Load balancing algorithm. This parameter is valid only when <b>readMode</b> is set to <b>SLAVE</b> or <b>MASTER_SLAVE</b> . Read traffic is distributed evenly.
masterConnecti onMinimumIdle Size	32	Minimum number of connections to the master node of each shard.
masterConnecti onPoolSize	64	Maximum number of connections to the master node of each shard.
slaveConnection MinimumIdleSiz e	32	Minimum number of connections to each replica node of each shard. If <b>readMode</b> is set to <b>MASTER</b> , the value of this parameter is invalid.
slaveConnection PoolSize	64	Maximum number of connections to each replica node of each shard. If <b>readMode</b> is set to <b>MASTER</b> , the value of this parameter is invalid.
subscriptionMod e	SLAVE	Subscription mode. By default, only replica nodes handle subscription. The value can be <b>SLAVE</b> or <b>MASTER</b> (recommended).
subscriptionCon- nectionMinimu mIdleSize	1	Minimum number of connections to the target node for pub/sub.
subscriptionCon- nectionPoolSize	50	Maximum number of connections to the target node for pub/sub.
subcriptionPerC onnection	5	Maximum number of subscriptions on each subscription connection.
connectionTime out	10000	Connection timeout interval, in milliseconds.
idleConnectionTi meout	10000	Maximum time (in milliseconds) for reclaiming idle connections.

Parameter	Default Value	Description
pingConnectionI nterval	30000	Heartbeat for detecting available connections, in milliseconds. <b>Recommended: 3000 ms</b> .
timeout	3000	Timeout interval for waiting for a response, in milliseconds.
retryAttempts	3	Maximum number of retries upon send failures.
retryInterval	1500	Retry interval, in milliseconds. <b>Recommended: 200 ms</b> .
clientName	null	Client name.

 Table 4-18
 ClusterServersConfig
 parameters
 (Redis
 Cluster)

Parameter	Default Value	Description
nodeAddress	-	Connection addresses of cluster nodes. Each address uses the redis:// <i>ip:port</i> format. Use commas (,) to separate connection addresses of different nodes.
password	null	Password for logging in to the cluster.
scanInterval	1000	Interval for periodically checking the cluster node status, in milliseconds.
readMode	SLAVE	Read mode. By default, read traffic is distributed to replica nodes. The value can be <b>MASTER</b> (recommended), <b>SLAVE</b> , or <b>MASTER_SLAVE</b> . Other values may cause access failures in failover scenarios.
loadBalancer	RoundRobinLoa dBalancer	Load balancing algorithm. This parameter is valid only when <b>readMode</b> is set to <b>SLAVE</b> or <b>MASTER_SLAVE</b> . Read traffic is distributed evenly.
masterConnecti onMinimumIdle Size	32	Minimum number of connections to the master node of each shard.
masterConnecti onPoolSize	64	Maximum number of connections to the master node of each shard.
slaveConnection MinimumIdleSiz e	32	Minimum number of connections to each replica node of each shard. If <b>readMode</b> is set to <b>MASTER</b> , the value of this parameter is invalid.

Parameter	Default Value	Description
slaveConnection PoolSize	64	Maximum number of connections to each replica node of each shard. If <b>readMode</b> is set to <b>MASTER</b> , the value of this parameter is invalid.
subscriptionMod e	SLAVE	Subscription mode. By default, only replica nodes handle subscription. The value can be <b>SLAVE</b> or <b>MASTER</b> (recommended).
subscriptionCon- nectionMinimu mIdleSize	1	Minimum number of connections to the target node for pub/sub.
subscriptionCon- nectionPoolSize	50	Maximum number of connections to the target node for pub/sub.
subcriptionPerC onnection	5	Maximum number of subscriptions on each subscription connection.
connectionTime out	10000	Connection timeout interval, in milliseconds.
idleConnectionTi meout	10000	Maximum time (in milliseconds) for reclaiming idle connections.
pingConnectionI nterval	30000	Heartbeat for detecting available connections, in milliseconds. <b>Recommended: 3000</b> .
timeout	3000	Timeout interval for waiting for a response, in milliseconds.
retryAttempts	3	Maximum number of retries upon send failures.
retryInterval	1500	Retry interval, in milliseconds. <b>Recommended: 200</b> .
clientName	null	Client name.

# Suggestion for Configuring DCS Instances

## • readMode

**MASTER** is the recommended value, that is, the master node bears all read and write traffic. This is to avoid data inconsistency caused by master/replica synchronization latency. If the value is **SLAVE**, all read requests will trigger errors when replicas are faulty. If the value is **MASTER\_SLAVE**, some read requests will trigger errors. Read errors last for the period specified by **failedSlaveCheckInterval** (180s by default) until the faulty nodes are removed from the available node list.

If read traffic and write traffic need to be separated, you can use read/write splitting DCS instances. Proxy nodes are deployed in the middle to distribute

read and write traffic. When a replica node is faulty, traffic is automatically switched to the master node. The switchover does not interrupt service applications, and the fault detection time window is far shorter than Redisson's window.

#### subscriptionMode

Similar to **readMode**, **MASTER** is the recommended value.

• Connection pool configuration

**NOTE** 

The following calculation is applicable only to common service scenarios. You can customize it based on your service requirements.

There is no standard connection pool size. You can configure one based on your service traffic. The following formulas are for reference:

- Minimum number of connections = (QPS of a single node accessing Redis)/(1000 ms/Average time spent on a single command)
- Maximum number of connections = (QPS of a single node accessing Redis)/(1000 ms/Average time spent on a single command) x 150%

For example, if the QPS of a service application is about 10,000, each request needs to access Redis 10 times (that is, 100,000 accesses to Redis every second), and the service application has 10 hosts, the calculation is as follows:

QPS of a single node accessing Redis = 100,000/10 = 10,000

Average time spent on a single command = 20 ms (Redis takes 5 ms to 10 ms to process a single command under normal conditions. If network jitter occurs, it takes 15 ms to 20 ms.)

Minimum number of connections = 10,000/(1000 ms/20 ms) = 200Maximum number of connections =  $10,000/(1000 \text{ ms}/20 \text{ ms}) \times 150\% = 300$ 

• Retry configuration

Redisson supports retries. You can set the following parameters based on service requirements. Generally, configure three retries, and set the retry interval to about 200 ms.

- retryAttempts: number of retry times
- **retryInterval**: retry interval

#### **NOTE**

In Redisson, some APIs are implemented through LUA, and the performance is low. You are advised to use Jedis instead of Redisson.

# 4.3.5 Connecting to Redis on redis-py (Python)

This section describes how to access a Redis instance on redis-py. For more information about how to use other Redis clients, visit **the Redis official website**.

The following operations are based on an example of accessing a Redis instance on a client on an elastic cloud server (ECS).

#### **Notes and Constraints**

Use redis-py to connect to single-node, master/standby, and Proxy Cluster instances and redis-py-cluster to connect to Redis Cluster instances.
To access a Redis 7.0 instance, use a redis-py **4.3.0** or later client. **5.0.0** and later versions are recommended.

# Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- An ECS has been created. For details about how to create an ECS, see Purchasing a Custom ECS
- If the ECS runs the Linux OS, ensure that the Python compilation environment has been installed on the ECS.
- The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

# Procedure

**Step 1** View the IP address/domain name and port of the DCS Redis instance to be accessed.

For details, see Viewing and Modifying Basic Settings of a DCS Instance in .

**Step 2** Log in to the ECS.

The following uses CentOS as an example to describe how to access an instance using a Python client.

Step 3 Access the DCS Redis instance.

If the ECS OS does not provide Python, run the following **yum** command to install it:

yum install python

#### 

The Python version must be 3.6 or later. If the default Python version is earlier than 3.6, perform the following operations to change it:

- 1. Run the **rm -rf python** command to delete the Python symbolic link.
- 2. Run the **In -s python***X.X.X* **python** command to create another Python link. In the command, *X.X.X* indicates the Python version number.
- For single-node, master/standby, or Proxy Cluster types:
  - a. Install Python and redis-py.
    - i. If the system does not provide Python, run the **yum** command to install it.
    - Run the following command to download and decompress the redispy package: wget https://github.com/andymccurdy/redis-py/archive/master.zip unzip master.zip
    - iii. Go to the directory where the decompressed redis-py package is saved, and install redis-py. python setup.py install

After the installation, run the **python** command. redis-py have been successfully installed if the following command output is displayed:

#### Figure 4-7 Running the python command

- b. Use the redis-py client to connect to the instance. In the following steps, commands are executed in CLI mode. (Alternatively, write the commands into a Python script and then execute the script.)
  - i. Run the **python** command to enter the CLI mode. You have entered CLI mode if the following command output is displayed:

#### Figure 4-8 Entering the CLI mode

[root@ecs]2222293 redis-py-master]# python
Python 3.6.8 (default, Nov 16 2020, 16:55:22)
[GCC 4.8.5 20150623 (Red Hat 4.8.5-44)] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>> import redis

ii. Run the following command to access the chosen DCS Redis instance:

r = redis.StrictRedis(host='XXX.XXX.XXX.XXX', port=6379, password='\*\*\*\*\*');

*XXX.XXX.XXX.XXX* indicates the IP address/domain name of the DCS instance and **6379** is an example port number of the instance. For details about how to obtain the IP address/domain name and port, see **Step 1**. Change them as required. \*\*\*\*\*\* indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

You have successfully accessed the instance if the following command output is displayed. Enter commands to perform read and write operations on the database.

#### Figure 4-9 Redis connected successfully



- For the Redis Cluster type:
  - a. Install the redis-py-cluster client.
    - i. Download the released version. wget https://github.com/Grokzen/redis-py-cluster/releases/download/2.1.3/redis-pycluster-2.1.3.tar.gz
    - ii. Decompress the package. tar -xvf redis-py-cluster-2.1.3.tar.gz
    - Go to the directory where the decompressed redis-py-cluster package is saved, and install redis-py-cluster. python setup.py install
  - b. Access the DCS Redis instance by using redis-py-cluster.

In the following steps, commands are executed in CLI mode. (Alternatively, write the commands into a Python script and then execute the script.)

- i. Run the **python** command to enter the CLI mode.
- Run the following command to access the chosen DCS Redis instance. If the instance does not have a password, exclude password='\*\*\*\*\*\*' from the command.
   >>> from rediscluster import RedisCluster

>>> startup\_nodes = [{"host": "192.168.0.143", "port": "6379"},{"host": "192.168.0.144", "port": "6379"},{"host": "192.168.0.145", "port": "6379"},{"host": "192.168.0.146", "port": "6379"}]

>>> rc = RedisCluster(startup\_nodes=startup\_nodes, decode\_responses=True, password='\*\*\*\*\*')

```
>>> rc.set("foo", "bar")
True
>>> print(rc.get("foo"))
'bar'
```

----End

# 4.3.6 Connecting to Redis on go-redis (Go)

This section describes how to access a Redis instance on go-redis. For more information about how to use other Redis clients, visit **the Redis official website**.

The following operations are based on an example of accessing a Redis instance on a client on an elastic cloud server (ECS).

## **Notes and Constraints**

To access a Redis 7.0 instance, use a go-redis 9.2.0 or later client.

## Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- View the IP address/domain name and port of the DCS Redis instance to be accessed. For details, see Viewing and Modifying Basic Settings of a DCS Instance.
- An ECS has been created. For details about how to create an ECS, see **Purchasing a Custom ECS**
- The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

# Connecting to Redis on go-redis

**Step 1** Log in to the ECS.

A Windows ECS is used as an example.

- **Step 2** Install Visual Studio Community 2017 on the ECS.
- **Step 3** Start Visual Studio and create a project. The project name can be customized. In this example, the project name is set to **redisdemo**.
- **Step 4** Import the dependency package of go-redis and enter **go get github.com/goredis/redis** on the terminal.

#### Step 5 Write the following code:

```
package main
import (
   "fmt'
   "github.com/go-redis/redis"
func main() {
  // Single-node
  rdb := redis.NewClient(&redis.Options{
     Addr: "host:port",
Password: "*******", // no password set
              0, // use default DB
     DB:
  })
  val, err := rdb.Get("key").Result()
  if err != nil {
     if err == redis.Nil {
        fmt.Println("key does not exists")
        return
     }
     panic(err)
  }
  fmt.Println(val)
  //Cluster
   rdbCluster := redis.NewClusterClient(&redis.ClusterOptions{
     Addrs: []string{"host:port"},
Password: "*******",
  })
  val1, err1 := rdbCluster.Get("key").Result()
  if err1 != nil {
     if err == redis.Nil {
        fmt.Println("key does not exists")
        return
     }
     panic(err)
  3
  fmt.Println(val1)
}
```

**host:port** are the IP address/domain name and port of the DCS Redis instance. For details about how to obtain the IP address/domain name and port, see **Prerequisites**. Change them as required. \*\*\*\*\*\*\*\* indicates the password used to log in to the DCS Redis instance. This password is defined during DCS Redis instance creation.

**Step 6** Run the **go build -o test main.go** command to package the code into an executable file, for example, **test**.

#### 

To run the package in the Linux OS, set the following parameters before packaging:

set GOARCH=amd64

#### set GOOS=linux

Step 7 Run the ./test command to access the DCS instance.

----End

# 4.3.7 Connecting to Redis on hiredis (C++)

This section describes how to access a Redis instance on hiredis (C++). For more information about how to use other Redis clients, visit **the Redis official website**.

The following operations are based on an example of accessing a Redis instance on a client on an elastic cloud server (ECS).

# Notes and Constraints

The operations described in this section apply only to single-node, master/standby, and Proxy Cluster instances. To use C++ to connect to a Redis Cluster instance, see the C++ Redis client description.

To access a Redis 7.0 instance, use a hiredis **1.1.0-rc1** or later client. For example, valkey 7.2.5 and later are recommended.

# Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- An ECS has been created. For details about how to create an ECS, see **Purchasing a Custom ECS**
- If the ECS runs the Linux OS, ensure that the GCC compilation environment has been installed on the ECS.
- The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

# Connecting to Redis on hiredis

**Step 1** View the IP address/domain name and port of the DCS Redis instance to be accessed.

For details, see Viewing and Modifying Basic Settings of a DCS Instance.

Step 2 Log in to the ECS.

The following uses CentOS as an example to describe how to access an instance in C++.

**Step 3** Install GCC, Make, and hiredis.

If the system does not provide a compiling environment, run the following **yum** command to install the environment:

yum install gcc make

- **Step 4** Run the following command to download and decompress the hiredis package: wget https://github.com/redis/hiredis/archive/master.zip unzip master.zip
- **Step 5** Go to the directory where the decompressed hiredis package is saved, and compile and install hiredis.

make make install Step 6 Access the DCS instance by using hiredis.

The following describes connection and password authentication of hiredis. For more information on how to use hiredis, visit the Redis official website.

 Edit the sample code for connecting to a DCS instance, and then save the code and exit. vim connRedis.c

```
Example:
#include <stdio.h>
#include <stdlib.h>
#include <string.h>
#include <hiredis.h>
int main(int argc, char **argv) {
   unsigned int j;
   redisContext *conn;
   redisReply *reply;
   if (argc < 3) {
        printf("Usage: example {instance_ip_address} 6379 {password}\n");
        exit(0):
   }
   const char *hostname = argv[1];
   const int port = atoi(argv[2]);
   const char *password = argv[3];
   struct timeval timeout = { 1, 500000 }; // 1.5 seconds
   conn = redisConnectWithTimeout(hostname, port, timeout);
   if (conn == NULL || conn->err) {
     if (conn) {
        printf("Connection error: %s\n", conn->errstr);
        redisFree(conn);
     } else {
        printf("Connection error: can't allocate redis context\n");
     }
   exit(1);
   }
   /* AUTH */
   reply = redisCommand(conn, "AUTH %s", password);
   printf("AUTH: %s\n", reply->str);
   freeReplyObject(reply);
   /* Set */
   reply = redisCommand(conn,"SET %s %s", "welcome", "Hello, DCS for Redis!");
   printf("SET: %s\n", reply->str);
   freeReplyObject(reply);
   /* Get */
   reply = redisCommand(conn,"GET welcome");
   printf("GET welcome: %s\n", reply->str);
   freeReplyObject(reply);
   /* Disconnects and frees the context */
   redisFree(conn);
   return 0;
}
```

2. Run the following command to compile the code: gcc connRedis.c -o connRedis -I /usr/local/include/hiredis -lhiredis

If an error is reported, locate the directory where the **hiredis.h** file is saved and modify the compilation command.

After the compilation, an executable **connRedis** file is obtained.

3. Run the following command to access the chosen DCS Redis instance: ./connRedis {redis\_instance\_address} 6379 {password}

*{redis\_instance\_address}* indicates the IP address/domain name of DCS instance and **6379** is an example port number of DCS instance. For details

about how to obtain the IP address/domain name and port, see **Step 1**. Change them as required. *{password}* indicates the password used to log in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

You have successfully accessed the instance if the following command output is displayed:

AUTH: OK SET: OK GET welcome: Hello, DCS for Redis!

## NOTICE

If an error is reported, indicating that the hiredis library files cannot be found, run the following commands to copy related files to the system directories and add dynamic links:

mkdir /usr/lib/hiredis cp /usr/local/lib/libhiredis.so.0.13 /usr/lib/hiredis/ mkdir /usr/include/hiredis cp /usr/local/include/hiredis/hiredis.h /usr/include/hiredis/ echo '/usr/local/lib' >>;>>;/etc/ld.so.conf ldconfig

Replace the locations of the **so** and **.h** files with actual ones before running the commands.

#### ----End

# 4.3.8 Connecting to Redis on StackExchange.Redis (C#)

This section describes how to access a Redis instance on StackExchange.Redis. For more information about how to use other Redis clients, visit **the Redis official website**.

The following operations are based on an example of accessing a Redis instance on a client on an elastic cloud server (ECS).

# Notes and Constraints

If you use the StackExchange client to connect to a Proxy Cluster instance, the multi-DB function cannot be used.

To access a Redis 7.0 instance, use a hiredis **2.6.111** or later client. **2.7.0** and later versions are recommended.

# Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- An ECS has been created. For details about how to create an ECS, see Purchasing a Custom ECS
- If the ECS runs the Linux OS, ensure that the GCC compilation environment has been installed on the ECS.

 The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

# Connecting to Redis on StackExchange.Redis

**Step 1** View the IP address/domain name and port of the DCS Redis instance to be accessed.

For details, see Viewing and Modifying Basic Settings of a DCS Instance.

**Step 2** Log in to the ECS.

A Windows ECS is used as an example.

- **Step 3** Install Visual Studio Community 2017 on the ECS.
- Step 4 Start Visual Studio 2017 and create a project.

Set the project name to redisdemo.

Step 5 Install StackExchange.Redis by using the NuGet package manager of Visual Studio.

Access the NuGet package manager console according to **Figure 4-10**, and enter **Install-Package StackExchange.Redis** - *Version 2.2.79*. (The version number is optional).

#### Figure 4-10 Accessing the NuGet package manager console

- Microsoft Visual Studio (Adm	<u>ninistr</u>	ator)				
Project Build Debug Team	Tools	Test Analyze	Window	Help		
🛀 💾 🚰 🦻 - 🤆 - 🛛 Debug 🕞		Set Tools and Feature	es		2	🗉 📕 위 개 개 🖕
×	Ċ,	xtensions and Updat	tes			
	* <b>e</b> (	onnect to Database.				- 🗣 redisConn
]using System;	Ť≣ (	Connect to Server				
using StackExchange.Redis;	١	Veb Code Analysis			•	
namespace redisdemo		ode Snippets Manae	ger	Ctrl+K, Ctrl+B		
{ ] class Program		hoose Toolbox Item	IS			
{	1	NuGet Package Mana	ager		•	வ Package Manager Console
private static Configura		reate GUID				Manage NuGet Packages for Solution
<pre>//the lock for singleton private static readonly</pre>	E	rror Lookup				Package Manager Settings
//singleton	E	xternal Tools				
private static Connectio		Import and Export Settings				
public static Connection		ustomize				
{ if (redisConn == nul	₿ •	)ptions				

**Step 6** Write the following code, and use the String Set and Get methods to test the connection.

using System; using StackExchange.Redis; namespace redisdemo { class Program { // redis config private static ConfigurationOptions connDCS = ConfigurationOptions.Parse("{instance\_ip\_address}: {port},password=\*\*\*\*\*\*\*\*,connectTimeout=2000"); //the lock for singleton private static readonly object Locker = new object(); //singleton private static ConnectionMultiplexer redisConn;

```
//singleton
   public static ConnectionMultiplexer getRedisConn()
     if (redisConn == null)
     {
        lock (Locker)
        {
           if (redisConn == null || !redisConn.IsConnected)
              redisConn = ConnectionMultiplexer.Connect(connDCS);
           }
        }
     }
     return redisConn;
  }
  static void Main(string[] args)
  {
     redisConn = getRedisConn();
     var db = redisConn.GetDatabase();
     //set get
     string strKey = "Hello";
     string strValue = "DCS for Redis!";
     Console.WriteLine( strKey + ", " + db.StringGet(strKey));
     Console.ReadLine();
  }
}
```

*{instance\_ip\_address}* and *{port}* are the IP address/domain name and port number of the DCS Redis instance. For details about how to obtain the IP address/ domain name and port, see **Step 1**. Change them as required. *\*\*\*\*\*\*\*\** indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

**Step 7** Run the code. You have successfully accessed the instance if the following command output is displayed: Hello, DCS for Redis!

For more information about other commands of StackExchange.Redis, visit **StackExchange.Redis**.

----End

3

# 4.3.9 Connecting to Redis on phpredis (PHP)

This section describes how to connect to Redis on phpredis. For more information about how to use other Redis clients, visit **the Redis official website**.

The following operations are based on an example of accessing a Redis instance on a client on an elastic cloud server (ECS).

# **Notes and Constraints**

The operations described in this section apply only to single-node, master/standby, and Proxy Cluster instances. To use phpredis to connect to a Redis Cluster instance, see the **phpredis description**.

## Prerequisites

• A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.

- An ECS has been created. For details about how to create an ECS, see **Purchasing a Custom ECS**
- If the ECS runs the Linux OS, ensure that the GCC compilation environment has been installed on the ECS.
- The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

# Connecting to Redis on phpredis

**Step 1** View the IP address/domain name and port of the DCS Redis instance to be accessed.

For details, see Viewing and Modifying Basic Settings of a DCS Instance.

Step 2 Log in to the ECS.

The following uses CentOS as an example to describe how to access an instance through phpredis.

**Step 3** Install GCC-C++ and Make compilation components.

#### yum install gcc-c++ make

Step 4 Install the PHP development package and CLI tool.

Run the following **yum** command to install the PHP development package:

#### yum install php-devel php-common php-cli

After the installation is complete, run the following command to query the PHP version and check whether the installation is successful:

#### php --version

- **Step 5** Install the phpredis client.
  - 1. Download the source phpredis package.

#### wget http://pecl.php.net/get/redis-5.3.7.tgz

This version is used as an example. To download phpredis clients of other versions, visit the Redis or PHP official website.

2. Decompress the source phpredis package.

# tar -zxvf redis-5.3.7.tgz

cd redis-5.3.7

3. Command before compilation.

# phpize

4. Configure the **php-config** file.

#### ./configure --with-php-config=/usr/bin/php-config

The location of the file varies depending on the OS and PHP installation mode. You are advised to locate the directory where the file is saved before the configuration.

#### find / -name php-config

5. Compile and install the phpredis client.

#### make && make install

6. After the installation, add the **extension** configuration in the **php.ini** file to reference the Redis module.

#### vim /etc/php.ini

Add the following configuration:

extension = "/usr/lib64/php/modules/redis.so"

#### **NOTE**

The **redis.so** file may be saved in a different directory from **php.ini**. Run the following command to locate the directory:

#### find / -name php.ini

7. Save the configuration and exit. Then, run the following command to check whether the extension takes effect:

#### php -m |grep redis

If the command output contains **redis**, the phpredis client environment has been set up.

#### **Step 6** Access the DCS instance by using phpredis.

#### 1. Edit a redis.php file.

```
<?php
%redis_host = "{redis_instance_address}";
%redis_port = {port};
%user_pwd = "{password}";
%redis = new Redis();
if ($redis->connect($redis_host, $redis_port) == false) {
    die($redis->getLastError());
    }
    if ($redis->auth($user_pwd) == false) {
        die($redis->getLastError());
    }
    if ($redis->set("welcome", "Hello, DCS for Redis!") == false) {
        die($redis->getLastError());
    }
    value = $redis->get("welcome");
    echo $value;
    $redis->close();
}
```

*{redis\_instance\_address}* indicates the example IP address/domain name of the DCS instance and *{port}* indicates the port number of the DCS instance. For details about how to obtain the IP address/domain name and port, see **Step 1**. Change them as required. *{password}* indicates the password used to log in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation. If password-free access is enabled, shield the **if** statement for password authentication.

2. Run the **php redis.php** command to access the DCS instance.

```
----End
```

# 4.3.10 Connecting to Redis on predis (PHP)

This section describes how to connect to Redis on predis. For more information about how to use other Redis clients, visit **the Redis official website**.

The following operations are based on an example of accessing a Redis instance on a client on an elastic cloud server (ECS).

# Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- An ECS has been created. For details about how to create an ECS, see **Purchasing a Custom ECS**
- If the ECS runs the Linux OS, ensure that the PHP compilation environment has been installed on the ECS.
- The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

# Connecting to Redis on predis

**Step 1** View the IP address/domain name and port of the DCS Redis instance to be accessed.

For details, see Viewing and Modifying Basic Settings of a DCS Instance.

- **Step 2** Log in to the ECS.
- **Step 3** Install the PHP development package and CLI tool. Run the following **yum** command:

yum install php-devel php-common php-cli

**Step 4** After the installation is complete, check the version number to ensure that the installation is successful.

Step 5 Download the Predis package to the /usr/share/php directory.

- 1. Run the following command to download the Predis source file: wget https://github.com/predis/predis/archive/refs/tags/v2.2.2.tar.gz
  - **NOTE**

This version is used as an example. To download Predis clients of other versions, visit the Redis or PHP official website.

- 2. Run the following commands to decompress the source Predis package: tar -zxvf predis-2.2.2.tar.gz
- Rename the decompressed Predis directory predis and move it to /usr/ share/php/. mv predis-2.2.2 predis

- **Step 6** Edit a file used to connect to Redis.
  - Example of using **redis.php** to connect to a single-node, master/standby, or Proxy Cluster DCS Redis instance:

```
<?php
require 'predis/autoload.php';
Predis\Autoloader::register();
$client = new Predis\Client([
 'scheme' => 'tcp',
 'host' => '{redis_instance_address}',
 'port' =>{port},
 'password' => '{password}'
]);
$client->set('foo', 'bar');
$value = $client->qet('foo');
```

```
echo $value;
?>
```

 Example code for using redis-cluster.php to connect to Redis Cluster: <?php</li>

```
require 'predis/autoload.php';
    $servers = array(
        'tcp://{redis_instance_address}:{port}'
    );
    $options = array('cluster' => 'redis');
    $client = new Predis\Client($servers, $options);
    $client->set('foo', 'bar');
    $value = $client->get('foo');
    echo $value;
?>
```

*{redis\_instance\_address}* indicates the actual IP address/domain name of the DCS instance and *{port}* is the actual port number of DCS instance. For details about how to obtain the IP address/domain name and port, see **Step 1**. Change them as required. *{password}* indicates the password used to log in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation. If password-free access is required, delete the line that contains "password".

Step 7 Run the php redis.php command to access the DCS instance.

----End

# 4.3.11 Connecting to Redis on ioredis (Node.js)

This section describes how to access a Redis instance on ioredis. For more information about how to use other Redis clients, visit **the Redis official website**.

The following operations are based on an example of accessing a Redis instance on a client on an elastic cloud server (ECS).

## Notes and Constraints

The operations described in this section apply only to single-node, master/standby, and Proxy Cluster instances. To access a Redis Cluster instance on ioredis, see **Node.js Redis client description**.

# Prerequisites

- A Redis instance is created, and is in the **Running** state. To create a Redis instance, see **Buying a DCS Redis Instance**.
- An ECS has been created. For details about how to create an ECS, see Purchasing a Custom ECS
- If the ECS runs the Linux OS, ensure that the GCC compilation environment has been installed on the ECS.
- The client and the Redis instance must be interconnected before connecting to the instance. For details, see Network Conditions for Accessing DCS Redis.

# Connecting to Redis on ioredis

- For client servers running Ubuntu (Debian series):
- **Step 1** View the IP address/domain name and port of the DCS Redis instance to be accessed.

#### For details, see Viewing and Modifying Basic Settings of a DCS Instance.

- Step 2 Log in to the ECS.
- Step 3 Install Node.js.

apt install nodejs-legacy

If the preceding command does not work, run the following commands:

wget https://nodejs.org/dist/v0.12.4/node-v0.12.4.tar.gz --no-check-certificate tar -xvf node-v4.28.5.tar.gz cd node-v4.28.5 ./configure make make install

#### **NOTE**

After the installation is complete, run the **node --version** command to query the Node.js version to check whether the installation is successful.

- **Step 4** Install the node package manager (npm).
- **Step 5** Install the Redis client ioredis.
- **Step 6** Edit the sample script for connecting to a DCS Redis instance.

Add the following content to the **ioredisdemo.js** script, including information about connection and data reading.

```
var Redis = require('ioredis');
var redis = new Redis({
                    // Redis port
 port: 6379,
 host: '192.168.0.196', // Redis host
 tamily: 4, // 4 (IPv4) or 6 (IPv6)
 db: 0
});
redis.set('foo', 'bar');
redis.get('foo', function (err, result) {
 console.log(result);
}):
// Or using a promise if the last argument isn't a function
redis.get('foo').then(function (result) {
 console.log(result);
}):
// Arguments to commands are flattened, so the following are the same:
redis.sadd('set', 1, 3, 5, 7);
redis.sadd('set', [1, 3, 5, 7]);
// All arguments are passed directly to the redis server:
redis.set('key', 100, 'EX', 10);
```

**host** indicates the example IP address/domain name of the DCS instance and **port** indicates the port of the DCS instance. For details about how to obtain the IP address/domain name and port, see **Step 1**. Change them as required. \*\*\*\*\*\* indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

**Step 7** Run the sample script to access the chosen DCS Redis instance. node ioredisdemo.js

----End

#### • For client servers running CentOS (Red Hat series):

**Step 1** View the IP address/domain name and port of the DCS Redis instance to be accessed.

For details, see Viewing and Modifying Basic Settings of a DCS Instance.

- **Step 2** Log in to the ECS.
- Step 3 Install Node.js.
  - yum install nodejs

If the preceding command does not work, run the following commands:

```
wget https://nodejs.org/dist/v0.12.4/node-v0.12.4.tar.gz --no-check-certificate
tar -xvf node-v0.12.4.tar.gz
cd node-v0.12.4
./configure
make
make install
```

#### **NOTE**

After the installation is complete, run the **node --version** command to query the Node.js version to check whether the installation is successful.

Step 4 Install npm.

yum install npm

**Step 5** Install the Redis client ioredis.

npm install ioredis

Step 6 Edit the sample script for connecting to a DCS Redis instance.

Add the following content to the **ioredisdemo.js** script, including information about connection and data reading.

```
var Redis = require('ioredis');
var redis = new Redis({
 port: 6379,
                 // Redis port
 host: '192.168.0.196', // Redis host
 family: 4, // 4 (IPv4) or 6 (IPv6) password: '*****',
 db: 0
});
redis.set('foo', 'bar');
redis.get('foo', function (err, result) {
 console.log(result);
});
// Or using a promise if the last argument isn't a function
redis.get('foo').then(function (result) {
 console.log(result);
});
// Arguments to commands are flattened, so the following are the same:
redis.sadd('set', 1, 3, 5, 7);
redis.sadd('set', [1, 3, 5, 7]);
// All arguments are passed directly to the redis server:
redis.set('key', 100, 'EX', 10);
```

**host** indicates the example IP address/domain name of the DCS instance and **port** indicates the port of the DCS instance. For details about how to obtain the IP address/domain name and port, see **Step 1**. Change them as required. \*\*\*\*\*\* indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

**Step 7** Run the sample script to access the chosen DCS Redis instance. node ioredisdemo.js

----End

# 4.4 Connecting to Redis on the Console

Access a DCS Redis instance through Web CLI.

## **Notes and Constraints**

- The instance is in the **Running** state.
- Available only for DCS Redis 4.0 or later instances.
- Some commands cannot be run on Web CLI. For details, see Web CLI Commands.
- Do not enter sensitive information in Web CLI to avoid disclosure.
- If the value is empty, **nil** is returned after the **GET** command is executed.

# Connecting to Redis on the Console

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

- **Step 2** Click <sup>(2)</sup> in the upper left corner of the console and select the region where your instance is located.
- Step 3 In the navigation pane, choose Cache Manager. In the Operation column of the instance, choose More > Connect to Redis, as shown in the following figure.

Figure 4-11 Accessing Web CLI

Q Select a property or enter a keyword.											Q (\$
✓ Name ⇔	Status 🖯	Cache Engine \ominus	Type \ominus	CPU 🖯	Specifi	Used/Av	ailable Mem \ominus	Connection \ominus	Enterprise Pr	Billing Mode	Operation
dcs-up99 5da80b58-0e22-4dab-92f0-ef13bdc	Running	Redis 5.0	Master/Standby	x86		1	- 2/1,024 (0.2%)	redis-5da80b	default	Pay-per-Use Created on Sep 2	View Metric Restart More A
Total Records: 1 10 V < 1 >											Modify Specifications Connect to Redis
											Change Password Reset Password Master/Standby Switchover Clear Data Command Renaming Change Billing
											Delete

**Step 4** Enter the access password of the DCS instance. On Web CLI, select the current Redis database, enter a Redis command in the command box, and press **Enter**.

**NOTE** 

- If no operation is performed for more than 15 minutes, the connection times out. You must enter the access password to connect to the instance again.
- You do not need to enter a password for accessing a password-free DCS Redis instance.

----End

# 4.5 Public Access to a DCS Redis 3.0 Instance (Discontinued)

# 4.5.1 Enabling Public Access of a DCS Redis 3.0 Instance

If public access has been enabled for the instance, skip this section.

If public access is not enabled, follow the instructions in this section. You can enable or disable SSL encryption when enabling public access.

#### **NOTE**

- Before accessing a DCS instance through a public network (with SSL encryption), download a CA certificate to verify the certificate of the instance for security purposes.
- When accessing a DCS instance through a public network (without SSL encryption), access the EIP and port 6379 of the instance. You do not need to download certificates or install Stunnel on your client.
- You are advised to enable SSL to encrypt the data transmitted between your Redis client and DCS instance to prevent data leakage.

# Prerequisites

- Cache engine: Must be Redis 3.0. Otherwise, public access cannot be enabled.
- Password protected: Must be yes. If not, enable password protection for the instance by referring to **Resetting an Instance Password**.

# **Enabling Public Access of a DCS Redis Instance**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of the DCS Redis instance you want to configure. A page with basic information about the DCS instance is displayed.
- **Step 5** Click on the right of **Public Access**.
- **Step 6** Click **O** to enable public access.
- Step 7 Select an EIP from the Elastic IP Address drop-down list.

If no EIPs are available, click **View Elastic IP** to create an EIP on the network console. After an EIP is created, click the refresh button on the right of **Elastic IP Address** to select the new EIP.

**Step 8** (Optional) Enable or disable SSL as required.

You are advised to enable SSL to encrypt the data transmitted between your Redis client and DCS instance to prevent data leakage.

#### Step 9 Click OK.

It takes 1 to 2 minutes to enable public access.

You will be automatically redirected to the **Background Tasks** page, where the progress of the current task is displayed. If the task status is **Succeeded**, public access is successfully enabled.

----End

# 4.5.2 Connecting to Redis 3.0 over a Public Network on rediscli

This section describes how to access a Redis 3.0 instance over a public network on redis-cli.

Public access helps R&D personnel establish local environment for development or testing, improving development efficiency. However, in the production environment (official environment), access a DCS Redis instance through a VPC to ensure efficient access.

# Prerequisites

Before using redis-cli to access a DCS Redis instance over a public network, ensure that:

- The instance version is Redis 3.0 and public access has been enabled.
- If certificates are required for accessing the DCS instance, download the certificate from the DCS instance details page.

# Accessing Redis 3.0 over a Public Network (Linux and SSL Enabled)

**Step 1** Ensure that the security group rule allows public access through port 36379.

When SSL encryption is enabled, allow public access through port 36379. Ensure that the Stunnel client has been installed.

< sg-DCS		
Summary Inbound Rules Outbound Rules Associated Insta	inces	
Add Rule Fast-Add Rule Delete Allow Common Ports	nbound Rules: 7 Learn more	about security group configuration.
Protocol & Port 🍸 🕐	Туре	Source (?)
All	IPv4	sg-DCS 🕐
icmp : All	IPv4	0.0.0.0/0 (?)
TCP: 22	IPv4	0.0.0.0/0 ⑦
TCP: 80	IPv4	0.0.0.0/0 (?)
TCP : 443	IPv4	0.0.0.0/0 (?)
TCP : 3389	IPv4	0.0.0.0/0 (?)
TCP : 36379	IPv4	192. 100. 64/32

Figure 4-12 Security group rule (port 36379)

- **Step 2** Obtain the public access address and the certificates of the instance on the instance **Basic Information** page.
  - The public access address is displayed in the Connection section.
  - The certificates can be downloaded by clicking Download Certificate for Public Access in the Connection section. After decompression, you will obtain dcs-ca.cer (the public key certificate in binary format) and dcs-cabundle.pem (the certificate file in text format).

Figure 4-13 Viewing the public access address (SSL enabled; port 36379)

Public Access	On 🖉   ?		
	Public Access Address	13 36379	
	SSL	On 上 Download Certificate for Public Access	

- **Step 3** Log in to the local Linux device.
- **Step 4** Install the Stunnel client.

Use either of the following methods to install Stunnel.

#### **NOTE**

Installation methods **apt** and **yum** are recommended. Any common Linux OSs should support at least one of these installation methods.

• apt-get method:

**apt-get** is used to manage DEB software packages and applicable to Debian OSs such as Ubuntu. Run the following command to install Stunnel:

#### apt install stunnel or apt-get install stunnel

If you cannot find Stunnel after running the command, run the **apt update** command to update the configuration and then install Stunnel again.

• **yum** method:

**yum** is used to manage RPM software packages and applicable to OSs such as Fedora, CentOS, and Red Hat. Run the following command to install Stunnel:

#### yum install stunnel

- **Step 5** Open the Stunnel configuration file **stunnel.conf**.
  - If Stunnel is installed using **apt-get**, the configuration file is stored at the **/etc/stunnel/stunnel.conf** directory by default.

If this directory does not exist or no configuration file exists in it, add a directory or configuration file.

• If Stunnel is installed using **yum**, the configuration file is stored at the **/usr/ local/stunnel/stunnel.conf** directory by default.

If this directory does not exist or no configuration file exists in it, add a directory or configuration file.

#### **NOTE**

- If you are not sure where to store the configuration file, enter the **stunnel** command after the installation to view the directory for storing the configuration file.
- The configuration file can be stored in any directory. Specify this configuration file when starting Stunnel.
- **Step 6** Add the following content to the configuration file **stunnel.conf**, and then save and exit.

```
debug = 4
output = /var/log/stunnel.log
sslVersion = all
[redis-client]
client = yes
accept = 8000
connect = {public access address}
CAfile = /etc/stunnel/dcs-ca.cer
```

Modify the following parameters as required and leave other parameters unchanged:

- client: indicates Stunnel. The fixed value is yes.
- CAfile: specifies a CA certificate, which is optional. If a CA certificate is required, download and decompress the certificate dcs-ca.cer as instructed in Step 2. If it is not required, delete this parameter.
- **accept**: specifies the user-defined listening port number of Stunnel. Specify this parameter when accessing a DCS instance by using a Redis client.
- **connect**: specifies the forwarding address and port number of Stunnel. Set this parameter to the instance public access address obtained in **Step 2**.

The following is a configuration example:

```
[redis-client]
client = yes
CAfile = D:\tmp\dcs\dcs-ca.cer
accept = 8000
connect = 49.**.**.211:36379
```

#### **Step 7** Run the following commands to start Stunnel:

#### stunnel /{customdir}/stunnel.conf

In the preceding command, {*customdir*} indicates the customized storage directory for the **stunnel.conf** file described in **Step 5**. The following is a command example:

#### stunnel /etc/stunnel/stunnel.conf

#### 

For the Ubuntu OS, run the **/etc/init.d/stunnel4 start** command to start Stunnel. The service or process name is **stunnel4** for the Stunnel 4.x version.

After starting the Stunnel client, run the **ps** -**ef**|**grep stunnel** command to check whether the process is running properly.

**Step 8** Run the following command to check whether Stunnel is being listened:

#### netstat -plunt |grep 8000|grep "LISTEN"

**8000** indicates the user-defined listening port number of Stunnel configured in the **accept** field in **Step 6**.

If a line containing the port number **8000** is displayed in the returned result, Stunnel is running properly. When the Redis client connects to the address **127.0.0.1:8000**, Stunnel will forward requests to the DCS Redis instance.

#### **Step 9** Access the DCS Redis instance.

- 1. Log in to the local Linux device.
- 2. Run the following command to download the source code package of your Redis client from http://download.redis.io/releases/redis-5.0.8.tar.gz:

#### wget http://download.redis.io/releases/redis-5.0.8.tar.gz

#### **NOTE**

You can also install the Redis client by running the following yum or apt command:

- yum install redis
- apt install redis-server
- 3. Run the following command to decompress the source code package of your Redis client:

#### tar -xzf redis-5.0.8.tar.gz

4. Run the following commands to go to the Redis directory and compile the source code of your Redis client:

cd redis-5.0.8

make

5. Run the following commands to access the chosen DCS Redis instance: **cd src** 

./redis-cli -h 127.0.0.1 -p 8000

# 

In the preceding command:

- The address following -h indicates the address of the Stunnel client, which is 127.0.0.1.
- The port following -p is the listening port of the Stunnel client, which has been configured in the accept field in Step 6. 8000 is used an example.

Do not use the public access address and port displayed on the console for the **-h** and **-p** parameters.

6. Enter the password. You can read and write cached data only after the password is verified.

#### auth {password}

*{password}* indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

You have successfully accessed the instance if the following command output is displayed:

OK 127.0.0.1:8000>

----End

## Accessing Redis 3.0 over a Public Network (Linux and SSL Disabled)

**Step 1** Ensure that the security group rule allows public access through port 6379.

When SSL encryption is disabled, the instance public access address can be accessed only if access through port 6379 is allowed.

#### Figure 4-14 Security group rule (port 6379)

< sg-DCS		
Summary Inbound Rules Outbound Rules Associated Instances		
Add Rule Fast-Add Rule Delete Allow Common Ports Inbound F	Rules: 7 Learn more about sec	curity group configuration.
Protocol & Port 🝸 ③	Туре	Source ⑦
All	IPv4	sg-DCS ?
ICMP : All	IPv4	0.0.0.0/0 ⑦
TCP : 22	IPv4	0.0.0.0/0 ⑦
TCP:80	IPv4	0.0.0.0/0 ⑦
TCP : 443	IPv4	0.0.0.0/0 ⑦
TCP: 3389	IPv4	0.0.0.0/0 ⑦
TCP : 6379	IPv4	192.1

Step 2 Obtain the public access address of the instance.

The public access address is displayed in the **Connection** section of the instance **Basic Information** page.

Figure 4-15 Viewing the public access address (SSL disabled; port 6379)

ic Access	On 🖉 💿	
	Public Access Address	13 6379
	SSL	Off

- **Step 3** Log in to the local Linux device.
- **Step 4** Run the following command to download the source code package of your Redis client from http://download.redis.io/releases/redis-5.0.8.tar.gz:

#### wget http://download.redis.io/releases/redis-5.0.8.tar.gz

#### **NOTE**

Publ

You can also install the Redis client by running the following yum or apt command:

- yum install redis
- apt install redis-server
- **Step 5** Run the following command to decompress the source code package of your Redis client:

tar -xzf redis-5.0.8.tar.gz

**Step 6** Run the following commands to go to the Redis directory and compile the source code of your Redis client:

cd redis-5.0.8

make

Step 7 Run the following commands to access the chosen DCS Redis instance:

cd src

./redis-cli -h {public access address} -p 6379

Replace {public access address} with the address obtained in Step 2. For example:

./redis-cli -h 49.\*\*.\*\*.211 -p 6379

**Step 8** Enter the password. You can read and write cached data only after the password is verified.

#### auth {password}

*{password}* indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

You have successfully accessed the instance if the following command output is displayed:

OK 49.\*\*.\*\*.211:6379>

----End

# Accessing Redis 3.0 over a Public Network (Windows and SSL Enabled)

**Step 1** Ensure that the security group rule allows public access through port 36379.

When SSL encryption is enabled, allow port 36379 for public access. In this case, ensure that the Stunnel client has been installed.

Figure 4-16 Security group rule (port 36379)

< sg-DCS					
Summary Inbound Rules Outbound Rules Associated In	stances				
Add Rule Fast-Add Rule Delete Allow Common Ports	Inbound Rules: 7 Learn more ab	out security group configuration.			
Protocol & Port 7 ?	Туре	Source (?)			
All	IPv4	sg-DCS 🕐			
ICMP : All	IPv4	0.0.0.0/0 (?)			
TCP:22	IPv4	0.0.0.0/0 (?)			
TCP : 80	IPv4	0.0.0.0/0 (?)			
TCP : 443	IPv4	0.0.0.0/0 (?)			
TCP: 3389	IPv4	0.0.0.0/0 (?)			
TCP : 36379	IPv4	192			

**Step 2** Obtain the public access address and the certificates of the instance.

- The public access address is displayed in the **Connection** section.
- The certificates can be downloaded by clicking Download Certificate for Public Access in the Connection section. After decompression, you will obtain dcs-ca.cer (the public key certificate in binary format) and dcs-cabundle.pem (the certificate file in text format).

Figure 4-17 Viewing the public access address (SSL enabled; port 36379)

Public Access	On 🖉	
	Public Access Address	13 36379
	SSL	On 🛃 Download Certificate for Public Access

- Step 3 Download the latest Windows Stunnel installation package (for example, stunnel-5.44-win32-installer.exe) from https://www.stunnel.org/ downloads.html to the local Windows device.
- **Step 4** Run the Stunnel installation program and install the Stunnel client.
- Step 5 Configure the Stunnel client: Right-click on the taskbar and choose Edit Configuration. Add the following configuration and then save and exit. [redis-client] client = yes

```
CAfile = D:\tmp\dcs\dcs-ca.cer
accept = 8000
connect = {public access address}
```

Modify the following parameters as required and leave other parameters unchanged:

- client: indicates Stunnel. The fixed value is yes.
- CAfile: specifies a CA certificate, which is optional. If a CA certificate is required, download and decompress the certificate dcs-ca.cer as instructed in Step 2. If it is not required, delete this parameter.
- **accept**: specifies the user-defined listening port number of Stunnel. Specify this parameter when accessing an instance on a Redis client.
- **connect**: specifies the service address and port of Stunnel. Set this parameter to the instance public access address obtained in **Step 2**.

When SSL encryption is enabled, the configuration is similar to the following:

```
[redis-client]
client = yes
CAfile = D:\tmp\dcs\dcs-ca.cer
accept = 8000
connect = 49.**.**.211:36379
```

- **Step 6** Right-click **O** on the taskbar and choose **Reload Configuration**.
- **Step 7** Open the CLI tool **cmd.exe** and run the following command to check whether 127.0.0.1:8000 is being listened:

netstat -an |find "8000"

Assume that port **8000** is configured as the listening port on the client.

If **127.0.0.1:8000** is displayed in the returned result and its status is **LISTENING**, the Stunnel client is running properly. When the Redis client connects to the address **127.0.0.1:8000**, Stunnel will forward requests to the DCS Redis instance.

- **Step 8** Access the DCS Redis instance.
  - 1. Obtain and decompress the Redis client installation package.

The Windows Redis client installation package can be downloaded here.

2. Open the CLI tool **cmd.exe** and run commands to go to the directory where the decompressed Redis client installation package is saved.

For example, to go to the **D:\redis-64.3.0.503** directory, run the following commands:

D:

#### cd D:\redis-64.3.0.503

 Run the following commands to access the chosen DCS Redis instance: redis-cli -h 127.0.0.1 -p 8000 -a password>

# 

In the preceding command: The address following **-h** indicates the address of the Stunnel client, which is **127.0.0.1**. The port following **-p** is the listening port of the Stunnel client, which has been configured in the **accept** field in **Step 5**. **8000** is used an example. Do not use the public access address and port displayed on the console for the **-h** and **-p** parameters.

<password> indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

You have successfully accessed the instance if the following command output is displayed:

127.0.0.1:8000>

Enter **info** and the DCS instance information will be returned. If no information is returned or the connection is interrupted, right-click the Stunnel icon on the taskbar and choose **Show Log Window** from the shortcut menu to show logs of Stunnel for cause analysis.

----End

# Accessing Redis 3.0 over a Public Network (Windows and SSL Disabled)

**Step 1** Ensure that the security group rule allows public access through port 6379.

When SSL encryption is disabled, allow port 6379 for external access.

< sg-DCS				
Summary Inbound Rules Outbound Rules Associated In	stances			
Add Rule Fast-Add Rule Delete Allow Common Ports	Inbound Rules: 7 Learn more abo	out security group configuration.		
Protocol & Port 🍞 🕐	Туре	Source 🕐		
All	IPv4	sg-DCS ?		
ICMP : All	IPv4	0.0.0.0/0 ⑦		
TCP: 22	IPv4	0.0.0.0/0 (?)		
TCP: 80	IPv4	0.0.0.0/0 (?)		
TCP : 443	IPv4	0.0.0.0/0 (?)		
TCP: 3389	IPv4	0.0.0.0/0 (?)		
TCP: 6379	IPv4	192.101.64/32		

#### Figure 4-18 Security group rule (port 6379)

**Step 2** Obtain the public access address of the instance.

The public access address is displayed in the **Connection** section of the instance **Basic Information** page.

Figure 4-19 Viewing the public access address (SSL disabled; port 6379)

Public Access	On 🖉 🕐			
	Public Access Address	13 6379		
	SSL	Off		

**Step 3** Obtain and decompress the Redis client installation package.

The Windows Redis client installation package can be downloaded here.

**Step 4** Open the CLI tool **cmd.exe** and run commands to go to the directory where the decompressed Redis client installation package is saved.

For example, to go to the **D:\redis-64.3.0.503** directory, run the following commands:

D:

#### cd D:\redis-64.3.0.503

**Step 5** Run the following commands to access the chosen DCS Redis instance:

#### redis-cli -h {public network access IP} -p 6379 -a password>

In this command, *{public network access IP}* indicates the IP address of the DCS Redis instance obtained in **Step 2**. *<password>* indicates the password used for logging in to the chosen DCS Redis instance. This password is defined during DCS Redis instance creation.

You have successfully accessed the instance if the following command output is displayed:

139.\*\*.\*\*.175:6379>

Enter info and the DCS instance information will be returned.

----End

## Troubleshooting

• Symptom: "Error: Connection reset by peer" is displayed.

**Possible cause**: The security group is incorrectly configured. You need to enable port **36379** or **6379**.

• When redis-cli is used to connect to an instance, the following message is displayed indicating that the remote host forcibly closes an existing connection.

**Possible cause**: SSL encryption has been enabled, but Stunnel is not configured during connection. The IP address displayed on the console was used for connection. In this case, strictly follow the instructions provided in **Accessing Redis 3.0 over a Public Network (Linux and SSL Enabled)**.

• For more information about Redis connection failures, see **Troubleshooting Redis Connection Failures**.

# **5** Accessing a DCS Memcached Instance (Discontinued)

# 5.1 Configuring a Memcached Password

# Scenario

DCS Memcached instances can be accessed with or without passwords. After an instance is created, you can change its password setting in the following scenarios:

• If you want to access a password-protected DCS Memcached instance without the username and password, you can enable password-free access to clear the username and password of the instance.

The Memcached text protocol does not support username and password authentication. To access a DCS Memcached instance by using the Memcached text protocol, you must enable password-free access to the instance.

• If you want to access a password-free DCS Memcached instance using a username and password, you can set a password for the instance using the password reset function.

# Procedure

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

- **Step 2** Click O in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- Step 4 To change the password setting for a DCS Memcached instance, choose More > Reset Password in the Operation column in the row containing the chosen instance.
- **Step 5** In the **Reset Password** dialogue box, perform either of the following operations as required:

- From password-protected to password-free: Switch the toggle for **Password-Free Access** and click **OK**.
- From password-free to password-protected: Enter a password, confirm the password, and click **OK**.

----End

# 5.2 Connecting to Memcached on a Client

# 5.2.1 Connecting to Memcached on the Telnet

Access a DCS Memcached instance using a telnet client on an ECS in the same VPC.

# Prerequisites

- The DCS Memcached instance you want to access is in the Running state.
- An ECS has been created on which the client has been installed. For details on how to create ECSs, see the *Elastic Cloud Server User Guide*.

## **NOTE**

An ECS can communicate with a DCS instance that belongs to the same VPC and is configured with the same security group.

- If the ECS and DCS instance are in different VPCs, establish a VPC peering connection to achieve network connectivity between the ECS and DCS instance. For details, see Does DCS Support Cross-VPC Access?
- If different security groups have been configured for the ECS and DCS instance, set security group rules to achieve network connectivity between the ECS and DCS instance. For details, see **How Do I Configure a Security Group?**
- All annotations in the example code have been deleted.
- All command lines and code blocks are UTF-8 encoded. Using another encoding scheme will cause compilation problems or even command failures.

# Connecting to Memcached on the Telnet

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- Step 3 In the navigation pane, choose Cache Manager.
- **Step 4** On the **Cache Manager** page, click the name of the DCS Memcached instance you want to access. Obtain the IP address or domain name and port number of the instance.
- Step 5 Access the chosen DCS Memcached instance.
  - 1. Log in to the ECS.
  - 2. Run the following command to check whether telnet is installed on the ECS: which telnet

If the directory in which the telnet is installed is displayed, telnet has been installed on the ECS. If the client installation directory is not displayed, install the telnet manually.

D NOTE

- If telnet has not been installed in Linux, run the **yum -y install telnet** command to install it.
- In the Windows OS, choose Start > Control Panel > Programs > Programs and Features > Turn Windows features on or off, and enable telnet.
- 3. Run the following command to access the chosen DCS Memcached instance:

#### telnet {ip or domain name} {port}

In this command: *{ip or domain name}* indicates the IP address or domain name of the DCS Memcached instance. *{port}* indicates the port number of the DCS Memcached instance. Both the IP address or domain name and the port number are obtained in **Step 4**.

When you have successfully accessed the chosen DCS Memcached instance, information similar to the following is displayed:

Trying XXX.XXX.XXX.XXX... Connected to XXX.XXX.XXX.XXX. Escape character is '^]'.

#### D NOTE

- If **Password Protected** is not enabled for the instance, run the following commands directly after the instance is accessed successfully.
- If Password Protected is enabled for the instance, attempts to perform operations on the instance will result in the message "ERROR authentication required", indicating that you do not have the required permissions. In this case, enter auth username@password to authenticate first. username and password are that used for accessing the DCS Memcached instance.

Example commands for using the DCS Memcached instance (lines in bold are the commands and the other lines are the command output):

set hello 0 0 6 world! STORED get hello VALUE hello 0 6 world! END

----End

# 5.2.2 Connecting to Memcached on the Spymemcached (Java)

Access a DCS Memcached instance using a Java client on an ECS in the same VPC.

# Prerequisites

- The DCS Memcached instance you want to access is in the **Running** state.
- An ECS has been created on which the client has been installed. For details on how to create ECSs, see the *Elastic Cloud Server User Guide*.

#### D NOTE

An ECS can communicate with a DCS instance that belongs to the same VPC and is configured with the same security group.

- If the ECS and DCS instance are in different VPCs, establish a VPC peering connection to achieve network connectivity between the ECS and DCS instance. For details, see Does DCS Support Cross-VPC Access?
- If different security groups have been configured for the ECS and DCS instance, set security group rules to achieve network connectivity between the ECS and DCS instance. For details, see How Do I Configure a Security Group?
- The Java development kit (JDK) and common integrated development environments (IDEs) such as Eclipse have been installed on the ECS.
- You have obtained the **spymemcached**-*x*.*y*.*z*.**jar** dependency package.

#### **NOTE**

*x.y.z* indicates the version of the dependency package. The latest version is recommended.

# Connecting to Memcached on the Spymemcached

- Step 1 Log in to the DCS console.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** On the **Cache Manager** page, click the name of the DCS Memcached instance you want to access. Obtain the IP address or domain name and port number of the instance.
- **Step 5** Upload the obtained **spymemcached**-*x.y.z.***jar** dependency package to the created ECS.
- **Step 6** Log in to the ECS.
- **Step 7** Create a Java project on Eclipse and import the **spymemcached**-*x*.*y*.*z***.jar** dependency package. The project name is customizable.
- **Step 8** Create a **ConnectMemcached1** class, copy the following Java code to the class, and modify the code.
  - Example code for the password mode

Change *ip or domain name:port* to the IP address/domain name and port number obtained in **Step 4**. Set *userName* and *password* respectively to the username and password of the Memcached instance.

//Connect to the encrypted Memcached code using Java. import java.io.IOException; import java.util.concurrent.ExecutionException;

import net.spy.memcached.AddrUtil; import net.spy.memcached.ConnectionFactoryBuilder; import net.spy.memcached.ConnectionFactoryBuilder.Protocol; import net.spy.memcached.MemcachedClient; import net.spy.memcached.auth.AuthDescriptor; import net.spy.memcached.auth.PlainCallbackHandler; import net.spy.memcached.internal.OperationFuture; public class ConnectMemcached1 public static void main(String[] args) final String connectionaddress = "ip or domain name:port"; final String username = "userName";//Indicates the username. final String password = "*password*";//Indicates the password. MemcachedClient client = null; try { AuthDescriptor authDescriptor = new AuthDescriptor(new String[] {"PLAIN"}, new PlainCallbackHandler(username, password)); client = new MemcachedClient( new ConnectionFactoryBuilder().setProtocol(Protocol.BINARY) .setAuthDescriptor(authDescriptor) .build(), AddrUtil.getAddresses(connectionaddress)); String key = "memcached";//Stores data with the key being **memcached** in Memcached. String value = "Hello World";//The value is Hello World. int expireTime = 5; //Specifies the expiration time, measured in seconds. The countdown starts from the moment data is written. After the expireTime elapses, the data expires and can no longer be read. doExcute(client, key, value, expireTime);//Executes the operation. } catch (IOException e) ł e.printStackTrace(); } } \*Method of writing data to Memcached \*/ private static void doExcute(MemcachedClient client, String key, String value, int expireTime) ł try { OperationFuture<Boolean> future = client.set(key, expireTime, value); future.get();//spymemcached set () is asynchronous. future.get () waits until the cache.set () operation is completed, or does not need to wait. You can select based on actual requirements. System.out.println("The Set operation succeeded."); System.out.println("Get operation:" + client.get(key)); Thread.sleep(6000);//Waits for 6000 ms, that is, 6s. Then the data expires and can no longer be read. System.out.println("Perform the Get operation 6s later:" + client.get(key)); } catch (InterruptedException e) e.printStackTrace(); } catch (ExecutionException e) e.printStackTrace(); if (client != null) { client.shutdown(); 7 } }

• Example code for the password-free mode

Change *ip or domain name:port* to the IP address/domain name and port number obtained in **Step 4**.

//Connect to the password-free Memcached code using Java. import java.io.IOException; import java.util.concurrent.ExecutionException;

```
import net.spy.memcached.AddrUtil;
import net.spy.memcached.BinaryConnectionFactory;
import net.spy.memcached.MemcachedClient;
import net.spy.memcached.internal.OperationFuture;
public class ConnectMemcached
  public static void main(String[] args)
     final String connectionaddress = "ip or domain name:port";
     MemcachedClient client = null;
     try
     {
       client = new MemcachedClient(new BinaryConnectionFactory(),
AddrUtil.getAddresses(connectionaddress));
        String key = "memcached";//Stores data with the key being memcached in Memcached.
        String value = "Hello World";//The value is Hello World.
        int expireTime = 5; //Specifies the expiration time, measured in seconds. The countdown
starts from the moment data is written. After the expireTime elapses, the data expires and can no
longer be read.
        doExcute(client, key, value, expireTime);//Executes the operation.
     }
     catch (IOException e)
     {
        e.printStackTrace();
     }
  }
   *Method of writing data to Memcached
   */
  private static void doExcute(MemcachedClient client, String key, String value, int expireTime)
     try
     ł
        OperationFuture<Boolean> future = client.set(key, expireTime, value);
        future.get();//spymemcached set () is asynchronous. future.get () waits until the cache.set
() operation is completed, or does not need to wait. You can select based on actual requirements.
        System.out.println("The Set operation succeeded.");
        System.out.println("Get operation:" + client.get(key));
        Thread.sleep(6000);//Waits for 6000 ms, that is, 6s. Then the data expires and can no longer
be read.
        System.out.println("Perform the Get operation 6s later:" + client.get(key));
     }
     catch (InterruptedException e)
     {
        e.printStackTrace();
     }
     catch (ExecutionException e)
     ł
        e.printStackTrace();
     if (client != null)
     {
        client.shutdown();
     }
  }
}
```

**Step 9** Run the **main** method. The following result is displayed in the **Console** window of Eclipse:

The Set operation succeeded. Get operation: Hello World Perform the Get operation 6s later: null

----End

# 5.2.3 Connecting to Memcached on the Python-binarymemcached (Python)

Access a DCS Memcached instance using a Python client on an ECS in the same VPC.

# Prerequisites

- The DCS Memcached instance you want to access is in the **Running** state.
- Log in to the ECS. For details on how to create ECSs, see the *Elastic Cloud Server User Guide*.

#### D NOTE

An ECS can communicate with a DCS instance that belongs to the same VPC and is configured with the same security group.

- If the ECS and DCS instance are in different VPCs, establish a VPC peering connection to achieve network connectivity between the ECS and DCS instance. For details, see Does DCS Support Cross-VPC Access?
- If different security groups have been configured for the ECS and DCS instance, set security group rules to achieve network connectivity between the ECS and DCS instance. For details, see How Do I Configure a Security Group?
- Python has been installed on the ECS. The recommended version is 2.7.6 or later.
- You have obtained the **<u>python-binary-memcached-x.y.z.zip</u>** dependency package.

**NOTE** 

*x.y.z* indicates the version of the dependency package. The latest version is recommended.

# Connecting to Memcached on the Python-binary-memcached

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

- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- Step 3 In the navigation pane, choose Cache Manager.
- **Step 4** On the **Cache Manager** page, click the name of the DCS Memcached instance you want to access. Obtain the IP address or domain name and port number of the instance.
- **Step 5** Upload the obtained dependency package (for example, the **python-binary-memcached**-*x.y.z.***zip** package) to the created ECS.
- **Step 6** Log in to the ECS.
- **Step 7** Run the following commands to install the dependency package:

unzip -xzvf python-binary-memcached-x.y.z.zip

cd python-binary-memcached-x.y.z

python setup.py install

## D NOTE

If an error is reported during the installation, use the **apt** or **yum** installation method. For example, to install the dependency package by using the **apt** method, run the following commands:

## apt install python-pip;

#### pip install python-binary-memcached;

- **Step 8** Create a Python file named **dcs\_test.py**, copy the following Python code to the file, and modify the code.
  - Example code for the password mode

Change *ip or domain name:port* to the IP address/domain name and port number obtained in **Step 4**. Set *userName* and *password* respectively to the username and password of the Memcached instance. import bmemcached client = bmemcached.Client(('*ip or domain name:port*'), '*userName*', '*password*) print "set('key', 'hello world!')" print client.set('key', 'hello world!') print "get('key')" print client.get('key')

• Example code for the password-free mode

Change ip or domain name:port to the IP address/domain name and port number obtained in **Step 4**.

```
import bmemcached
client = bmemcached.Client('ip or domain name:port')
print "set('key', 'hello world!')"
print client.set('key', 'hello world!')
print "get('key')"
print client.get('key')
```

Step 9 Run the dcs\_test.py file. The following result is displayed.

# python test.py
set('key', 'hello world!')
True
get('key')
hello world!

----End

# 5.2.4 Connecting to Memcached on the Libmemcached (C++)

Access a DCS Memcached instance using a C++ client on an ECS in the same VPC.

# Prerequisites

- The DCS Memcached instance you want to access is in the **Running** state.
- Log in to the ECS. For details on how to create ECSs, see the *Elastic Cloud Server User Guide*.

#### D NOTE

An ECS can communicate with a DCS instance that belongs to the same VPC and is configured with the same security group.

- If the ECS and DCS instance are in different VPCs, establish a VPC peering connection to achieve network connectivity between the ECS and DCS instance. For details, see **Does DCS Support Cross-VPC Access**?
- If different security groups have been configured for the ECS and DCS instance, set security group rules to achieve network connectivity between the ECS and DCS instance. For details, see How Do I Configure a Security Group?
- GCC has been installed on the ECS. The recommended version is 4.8.4 or later.
- You have obtained the libmemcached-x.y.z.tar.gz dependency package.

#### **NOTE**

*x.y.z* indicates the version of the dependency package. The latest version is recommended.

# Connecting to Memcached on the Libmemcached (C++)

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** On the **Cache Manager** page, click the name of the DCS Memcached instance you want to access. Obtain the IP address or domain name and port number of the instance.
- **Step 5** Upload the obtained **libmemcached**-*x.y.z.***tar.gz** dependency package to the created ECS.
- **Step 6** Log in to the ECS.
- Step 7 Install related SASL dependency packages.

For OSs of Debian series: apt install libsasl2-dev cloog.ppl

For OSs of Red Hat series: yum install cyrus-sasl\*

**Step 8** Run the following commands to install the dependency package:

tar -xzvf libmemcached-x.y.z.tar.gz

cd libmemcached-x.y.z

./configure --enable-sasl

make

make install

**Step 9** Create a file named **build.sh** and copy the following code to the file. g++ -o dcs\_sample dcs\_sample.cpp -lmemcached -std=c++0x -lpthread -lsasl2
#### D NOTE

If the **libmemcached.so.11** file cannot be found during compilation, run the **find** command to find the file and copy the file to the **/usr/lib** directory.

- **Step 10** Create a file named **dcs\_sample.cpp**, copy the following C++ code to the file, and modify the code.
  - Example code for the password mode

```
Change ip or domain name and port to the IP address or domain name and
port obtained in Step 4. Set userName and password respectively to the
username and password of the Memcached instance.
#include <iostream>
#include <string>
#include <libmemcached/memcached.h>
using namespace std;
#define IP "ip or domain name"
#define PORT "port"
#define USERNAME "userName"
#define PASSWORD "password"
memcached_return result;
memcached_st * init()
ł
  memcached_st *memcached = NULL;
  memcached_server_st *cache;
  memcached = memcached_create(NULL);
  cache = memcached_server_list_append(NULL, IP, PORT, &result);
  sasl_client_init(NULL);
     memcached_set_sasl_auth_data(memcached, USERNAME, PASSWORD);
     memcached_behavior_set(memcached,MEMCACHED_BEHAVIOR_BINARY_PROTOCOL,1);
     memcached_server_push(memcached,cache);
     memcached_server_list_free(cache);
       return memcached;
}
int main(int argc, char *argv[])
{
     memcached_st *memcached=init();
  string key = "memcached";
  string value = "hello world!";
  size_t value_length = value.length();
  int expire_time = 0;
  uint32_t flag = 0;
  result =
memcached_set(memcached,key.c_str(),key.length(),value.c_str(),value.length(),expire_time,flag);
  if (result != MEMCACHED_SUCCESS){
   cout <<"set data failed: " << result << endl;
   return -1;
  }
  cout << "set succeed, key: " << key << ", value: " << value << endl;
  cout << "get key:" << key << endl;
  char* result = memcached_get(memcached,key.c_str(),key.length(),&value_length,&flag,&result);
  cout << "value:" << result << endl;
  memcached_free(memcached);
  return 0;
}
Example code for the password-free mode
Change ip and port to the IP address or domain name and port obtained in
Step 4.
```

```
#include <iostream>
#include <string>
```

```
#include <libmemcached/memcached.h>
using namespace std;
#define IP "ip or domain name"
#define PORT port
memcached_return result;
memcached_st * init()
{
  memcached_st *memcached = NULL;
  memcached_server_st *cache;
  memcached = memcached_create(NULL);
  cache = memcached_server_list_append(NULL, IP, PORT, &result);
      memcached_server_push(memcached,cache);
  memcached_server_list_free(cache);
      return memcached;
}
int main(int argc, char *argv[])
ł
      memcached_st *memcached=init();
  string key = "memcached";
  string value = "hello world!";
  size_t value_length = value.length();
  int expire_time = 0;
  uint32_t flag = 0;
  result =
memcached_set(memcached,key.c_str(),key.length(),value.c_str(),value.length(),expire_time,flag);
  if (result != MEMCACHED_SUCCESS){
    cout <<"set data failed: " << result << endl;</pre>
   return -1;
  }
  cout << "set succeed, key: " << key << ",value: " << value << endl;
  cout << "get key:" << key << endl;
  char* result = memcached_get(memcached,key.c_str(),key.length(),&value_length,&flag,&result);
  cout << "value:" << result << endl;
  memcached_free(memcached);
  return 0;
}
```

**Step 11** Run the following commands to compile the source code:

chmod 700 build.sh

./build.sh

The **dcs\_sample** binary file is generated.

Step 12 Run the following command to access the chosen DCS Memcached instance:

```
./dcs_sample
set succeed, key: memcached ,value: hello world!
get key:memcached
value:hello world!
```

```
----End
```

# 5.2.5 Connecting to Memcached on the Libmemcached (PHP)

Access a DCS Memcached instance using a PHP client on an ECS in the same VPC.

### Prerequisites

• The DCS Memcached instance you want to access is in the **Running** state.

• Log in to the ECS. For details on how to create ECSs, see the *Elastic Cloud Server User Guide*.

#### D NOTE

An ECS can communicate with a DCS instance that belongs to the same VPC and is configured with the same security group.

- If the ECS and DCS instance are in different VPCs, establish a VPC peering connection to achieve network connectivity between the ECS and DCS instance. For details, see Does DCS Support Cross-VPC Access?
- If different security groups have been configured for the ECS and DCS instance, set security group rules to achieve network connectivity between the ECS and DCS instance. For details, see **How Do I Configure a Security Group?**

# **OSs of Red Hat Series**

The following uses CentOS 7.0 as an example to describe how to install a PHP client and use it to access a DCS Memcached instance. The procedure is also applicable to a PHP client running the Red Hat or Fedora OS.

**Step 1** Install GCC-C++ and Make compilation components.

#### yum install gcc-c++ make

**Step 2** Install related SASL packages.

#### yum install cyrus-sasl\*

**Step 3** Install the libMemcached library.

Installing the libMemcached library requires SASL authentication parameters. Therefore, you cannot install the library by running the **yum** command.

wget https://launchpad.net/libmemcached/1.0/1.0.18/+download/ libmemcached-1.0.18.tar.gz

tar -xvf libmemcached-1.0.18.tar.gz

cd libmemcached-1.0.18

#### ./configure --prefix=/usr/local/libmemcached --enable-sasl

make && make install

**NOTE** 

Before installing the libMemcached library, install GCC-C++ and SASL components. Otherwise, an error will be reported during compilation. After you resolve the error, run the **make clean** command and then run the **make** command again.

#### **Step 4** Install the PHP environment.

yum install php-devel php-common php-cli

#### NOTICE

PHP 7.x does not support SASL authentication. Use PHP 5.6. If the yum php version is not 5.6, download one from the Internet.

#### **Step 5** Install the Memcached client.

Note that you must add a parameter used to enable SASL when running the **configure** command.

#### wget http://pecl.php.net/get/memcached-2.1.0.tgz

tar zxvf memcached-2.1.0.tgz

cd memcached-2.1.0

phpize

./configure --with-libmemcached-dir=/usr/local/libmemcached --enablememcached-sasl

#### make && make install

**Step 6** Modify the **php.ini** file.

Run the **find** or **locate** command to find the **php.ini** file.

#### find / -name php.ini

Add the following two lines to the **php.ini** file:

```
extension=memcached.so
memcached.use_sasl = 1
```

Figure 5-1 Modifying the php.ini file

#### Step 7 Access a DCS Memcached instance.

Create a **memcached.php** file and add the following content to the file:

php</th
<pre>\$connect = new Memcached; //Declares a Memcached connection.</pre>
<pre>\$connect-&gt;setOption(Memcached::OPT_COMPRESSION, false); //Disables compression.</pre>
<pre>\$connect-&gt;setOption(Memcached::OPT_BINARY_PROTOCOL, true); //Uses the binary protocol.</pre>
\$connect->setOption(Memcached::OPT_TCP_NODELAY, true); //Disables the TCP network delay policy.
\$connect->addServer('{memcached_instance_ip}, 11211); //Specifies the instance IP address and port.
<pre>\$connect-&gt;setSaslAuthData('{username}, '{password}); //If password-free access is enabled for the</pre>
instance, delete or comment out this line.
instance, delete or comment out this line.

```
$connect->set("DCS", "Come on!");
echo 'DCS: ',$connect->get("DCS");
echo "\n";
$connect->quit();
```

Save and run the **memcached.php** file. The following result is displayed.

[root@testphpmemcached ~]# php memcached.php DCS: Come on! [root@testphpmemcached ~]#

----End

### **OSs of Debian Series**

The following uses the Ubuntu OS as an example to describe how to install a PHP client and use it to access a DCS Memcached instance.

Step 1 Install GCC and Make compilation components.

#### apt install gcc make

**Step 2** Install the PHP environment.

PHP 5.x is recommended for better compatibility with SASL authentication.

Run the following commands to add the image source of PHP of an earlier version, and then install the **php.5.6** and **php.5.6-dev** packages:

#### apt-get install -y language-pack-en-base;

LC\_ALL=en\_US.UTF-8;

add-apt-repository ppa:ondrej/php;

apt-get update;

#### apt-get install php5.6 php5.6-dev;

After the installation is complete, run the **php** -version command to check the PHP version. If the following result is displayed, the PHP version is 5.6, indicating that PHP 5.6 is successfully installed.

root@dcs-nodelete:/etc/apt# php -version PHP 5.6.36-1+ubuntu16.04.1+deb.sury.org+1 (cli) Copyright (c) 1997-2016 The PHP Group

#### **NOTE**

To uninstall PHP, run the following commands:

apt install aptitude -y

aptitude purge `dpkg -l | grep php| awk '*{print \$2}*' |tr "\n" " "`

**Step 3** Install the SASL component.

#### apt install libsasl2-dev cloog.ppl

Step 4 Install the libMemcached library.

wget https://launchpad.net/libmemcached/1.0/1.0.18/+download/ libmemcached-1.0.18.tar.gz

tar -xvf libmemcached-1.0.18.tar.gz

#### cd libmemcached-1.0.18

#### ./configure --prefix=/usr/local/libmemcached

#### make && make install

**NOTE** 

Before installing the libMemcached library, install GCC-C++ and SASL components. Otherwise, an error will be reported during compilation. After you resolve the error, run the **make clean** command and then run the **make** command again.

**Step 5** Install the Memcached client.

Install the zlib component.

#### apt install zlib1g.dev

Note that you must add a parameter used to enable SASL when running the **configure** command.

wget http://pecl.php.net/get/memcached-2.2.0.tgz;

tar zxvf memcached-2.2.0.tgz;

cd memcached-2.2.0;

phpize5.6;

./configure --with-libmemcached-dir=/usr/local/libmemcached --enablememcached-sasl;

make && make install;

**Step 6** Modify the **pdo.ini** file.

Run the following command to find the **pdo.ini** file:

#### find / -name pdo.ini

By default, the **pdo.ini** file is stored in the **/etc/php/5.6/mods-available** directory. Add the following two lines to the **php.ini** file:

extension=memcached.so memcached.use\_sasl = 1

Figure 5-2 Modifying the pdo.ini file



**Step 7** Access a DCS Memcached instance.

Create a **memcached.php** file and add the following content to the file:

<?php

\$connect = new Memcached; //Declares a Memcached connection.
\$connect->setOption(Memcached::OPT\_COMPRESSION, false); //Disables compression.
\$connect->setOption(Memcached::OPT\_BINARY\_PROTOCOL, true); //Uses the binary protocol.

```
$connect->setOption(Memcached::OPT_TCP_NODELAY, true); //Disables the TCP network delay policy.
$connect->addServer('{memcached_instance_ip}", 11211); //Specifies the instance IP address and port.
$connect->setSaslAuthData('{username}", '{password}"); //If password-free access is enabled for the
instance, delete or comment out this line.
$connect->set("DCS", "Come on!");
echo 'DCS: ',$connect->get("DCS");
echo "\n";
$connect->quit();
?>
```

Save and run the **memcached.php** file. The following result is displayed.

[root@dcs-nodelete ~]# php memcached.php DCS: Come on! [root@dcs-nodelete ~]#

----End

# 6 Managing Instances

# 6.1 Viewing and Modifying Basic Settings of a DCS Instance

On the DCS console, you can view and modify DCS instance basic information.

# Viewing and Modifying Basic Information of a DCS Instance

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click O in the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Search for DCS instances using any of the following methods:
  - Search by keyword. Enter a keyword to search.
  - Select attributes and enter their keywords to search.

Currently, you can search by name, specification, ID, IP address, AZ, status, instance type, and cache engine.

For example, click the search box, choose **Cache Engine**, and then choose **Redis 3.0**, **Redis 4.0**, **Redis 5.0**, or **Redis 6.0**.

For more information on how to search, click the question mark to the right of the search box.

**Step 5** Click the name of the DCS instance to display more details about the DCS instance. **Table 6-1** describes the parameters.

| Section             | Parameter        | Description  |  |  |  |  |  |
|---------------------|------------------|--|--|--|--|--|--|
| Instance<br>Details | Name             | Name of the chosen instance. To modify the instance name, click 🧖.   |  |  |  |  |  |
|                     | Status           | State of the chosen instance.  |  |  |  |  |  |
|                     | ID               | ID of the chosen instance.   |  |  |  |  |  |
|                     | Cache<br>Engine  | Cache version of DCS. For example, Basic   Redis 4.0.<br>The cache version is fixed once the instance is<br>created. To use another version, create an instance<br>again and migrate the data.   |  |  |  |  |  |
|                     | Minor<br>Version | <ul> <li>Minor version of the instance. DCS optimizes functions and fixes vulnerabilities in minor upgrades. Clicking Upgrade to obtain the latest version.</li> <li>NOTE</li> <li>By default, a minor version cannot be viewed. To view or upgrade it, contact customer service to enable a whitelist.</li> <li>Perform instance upgrades during off-peak hours. Otherwise, upgrades may fail when the instance memory or CPU usage exceeds 90% or write traffic bursts. In such cases, try again during off-peak hours.</li> <li>The instance is upgraded by migrating nodes. During the migration, latency will increase. A migrating shard will become read-only for 1 minute and intermittently disconnected. Ensure that the client can reconnect and handle exceptions.</li> <li>For a Redis Cluster instance, ensure that the client can properly process the MOVED and ASK commands. Otherwise, requests will fail.</li> <li>This parameter is displayed only for basic edition DCS Redis 4.0 and later instances. Professional edition instances do not support minor version upgrades.</li> <li>If the minor version of the instance is already the latest, you cannot upgrade it.</li> <li>Upgrading the minor does not affect the instance connection addresses, password, whitelist, or</li> </ul> |  |  |  |  |  |

 Table 6-1 Parameters on the Basic Information page of a DCS instance

| Section | Parameter                    | Description   |
|---------|------------------------------|---|
|         | Proxy<br>Version             | Proxy version of a DCS instance. This parameter is<br>displayed only for Proxy Cluster and read/write<br>splitting instances. DCS optimizes functions and fixes<br>vulnerabilities in proxy upgrades. Clicking <b>Upgrade</b><br>to obtain the latest version.<br><b>NOTE</b> |
|         |                              | <ul> <li>By default, a proxy version cannot be viewed. To<br/>view or upgrade it, contact customer service to<br/>enable a whitelist.</li> </ul>  |
|         |                              | <ul> <li>The upgrading instance will be intermittently<br/>disconnected. Ensure that the client can reconnect and<br/>handle exceptions. Perform the upgrade during off-<br/>peak hours.</li> </ul>   |
|         |                              | <ul> <li>If the proxy version of the instance is already the latest,<br/>you cannot upgrade it.</li> </ul>  |
|         |                              | <ul> <li>Upgrading the proxy does not affect the instance<br/>connection addresses, password, whitelist, or<br/>monitoring and alarms.</li> </ul>   |
|         | Instance<br>Type             | Type of the selected instance. Currently, supported types include single-node, master/standby, Proxy Cluster, read/write splitting, and Redis Cluster.  |
|         |                              | To change the instance type, see <b>Modifying DCS</b><br><b>Instance Specifications</b> about the supported<br>instance types, changing notes and procedure.  |
|         | Cache Size                   | Specification of the chosen instance.   |
|         |                              | To change the instance specification, see <b>Modifying</b><br><b>DCS Instance Specifications</b> about the changing<br>notes and procedure.   |
|         | Bandwidth                    | Bandwidth of the DCS instance.  |
|         |                              | You can click <b>Adjust Bandwidth</b> to adjust the instance bandwidth. For details, see <b>Adjusting DCS Instance Bandwidth</b> .  |
|         | Used/<br>Available<br>Memory | The used memory space and maximum available memory space of the chosen instance.  |
|         | (MB)                         | The used memory space includes:   |
|         |                              | <ul> <li>Size of Redis-server buffers (including client<br/>buffer and repl-backlog) and internal data<br/>structures</li> </ul>  |
|         | CPU                          | CPU architecture of the chosen instance. This<br>parameter is displayed only for DCS Redis instances.<br>The CPU architecture is fixed once the instance is<br>created.   |

| Section        | Parameter             | Description  |
|----------------|-----------------------|--|
|                | Enterprise<br>Project | Enterprise project to which the new instance<br>belongs. Click is to modify the enterprise project of<br>the instance.<br>Enterprise projects isolate resources, personnel, and<br>finance. Modifying an enterprise project changes<br>isolated objectives.  |
|                | Maintenanc<br>e       | Time range for any scheduled maintenance activities<br>on cache nodes of this DCS instance. To modify the<br>window, click .<br>Select a new window from the drop-down list and<br>click ✓ to save, or × to cancel.<br>The modification takes effect immediately.  |
|                | Description           | Description of the chosen DCS instance. To modify the description, click 🧖.  |
| Connectio<br>n | Password<br>Protected | Yes: password-protected access; No: password-free<br>access.<br>To change the password access mode, see<br>Configuring a Redis Password.   |
|                | Connection<br>Address | The domain name and port of the Redis instance to be accessed on a client within the VPC.  |
|                |                       | <ul> <li>You can click next to Connection Address to change the port. The connection address is fixed once the instance is created.</li> <li>NOTE</li> <li>For a master/standby DCS Redis 4.0/5.0/6.0 instance, Connection Address indicates the domain name and port number of the master node, and Read-only Address indicates those of the standby node. When connecting to such an instance, you can use the domain name and port number of the master node or the standby node. For details, see the architecture of a</li> </ul> |
|                |                       | <ul> <li>master/standby instance.</li> <li>You can change the port only for a DCS Redis 4.0, 5.0, or 6.0 basic instance, but not for a DCS Redis 3.0, 6.0 professional, or Memcached instance.</li> </ul>  |
|                | IP Address            | The IP address and port of the DCS instance to be accessed on a client within the VPC.   |
|                |                       | To change the instance port, click . The IP address is fixed once the instance is created. The domain name address is recommended.   |

| Section              | Parameter         | Description   |
|----------------------|-------------------|---|
|                      | Public<br>Access  | Currently, public access can be enabled by default<br>only for Redis 3.0 instances. To enable public access<br>to a Redis 3.0 instance, see <b>Public Access to a DCS</b><br><b>Redis 3.0 Instance (Discontinued)</b> . Public access<br>cannot be enabled for instances of other versions.<br>To enable public access to a Redis 4.0, 5.0, or 6.0<br>instance, contact customer service to enable the<br>public access whitelist. <b>Then, see Enabling Public</b><br><b>Access to Redis and Obtaining the Access</b><br><b>Addresses to enable public access and obtain the</b><br><b>addresses</b> .<br>Memcached (discontinued) does not support public<br>access   |
| Network              | AZ                | Availability zone in which the instance nodes<br>running the selected DCS instance reside.<br>The AZ of standby nodes can be changed for a<br>cluster multi-replica instance in a single AZ. For<br>details, see <b>Changing Cluster DCS Instances to be</b><br><b>Across AZs</b> . AZs cannot be changed in other<br>scenarios.  |
|                      | VPC               | VPC in which the chosen instance resides. The VPC is fixed once the instance is created.  |
|                      | Subnet            | Subnet in which the chosen instance resides. The subnet is fixed once the instance is created.  |
|                      | Security<br>Group | Security group that controls access to the chosen<br>instance. To modify the security group, click .<br>Select a new security group from the drop-down list<br>and click to save, or to cancel. The<br>modification takes effect immediately.<br>Security group access control is supported only by<br>DCS Redis 3.0, professional edition DCS Redis 6.0,<br>and Memcached instances. Basic edition DCS for<br>Redis 4.0/5.0/6.0 is based on VPC Endpoint, does not<br>support security groups, and <b>Configuring DCS Redis</b><br><b>Access Whitelist</b> is suggested.<br><b>NOTE</b><br>Only the security groups that have been created can be<br>selected from the drop-down list. If you need to create a<br>security group, follow the procedure described in How Do I<br><b>Configure a Security Group?</b> |
| Instance<br>Topology | -                 | Hover over a node to view its metrics, or click the icon of a node to view its historical metrics.<br>Single-node instances do not display the instance   |
|                      |                   | topology.   |

| Section | Parameter       | Description   |
|---------|-----------------|---|
| Billing | Billing<br>Mode | Billing mode of the instance. To change the billing mode, see <b>Billing Mode Changes</b> . |
|         | Created         | Time at which the chosen instance started to be created.                                    |
|         | Run             | Time at which the instance was created.   |

----End

# 6.2 Viewing DCS Background Tasks

After you initiate certain instance operations such as scaling up the instance and changing or resetting a password, a background task will start for each operation. On the DCS console, you can view the background task status and clear task information by deleting task records.

# **Viewing DCS Background Tasks**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.

Filter DCS instances to find the desired DCS instance. Currently, you can search instances by name, specification, ID, IP address, AZ, status, instance type, cache engine, and many other attributes.

- **Step 4** Click the name of the DCS instance to display more details about the DCS instance.
- Step 5 Choose Background Tasks.

Filter tasks by specifying the time, property, or keyword.

- Click <sup>1</sup> to refresh the task status.
- To clear the record of a background task, choose **Operation** > **Delete**.

D NOTE

You can only delete the records of tasks in the **Successful** or **Failed** state.

----End

# 6.3 Viewing Client Connection Information of a DCS Instance

You can view the client connection information of a DCS instance and disconnect clients.

# **Notes and Constraints**

- The session management page displays only the information about the external client connections. Information about the Web CLI connections is not displayed.
- This function is available only in the CN North-Beijing1, CN North-Beijing4, CN East-Shanghai1, CN East-Shanghai2, CN South-Guangzhou, and CN Southwest-Guiyang1 regions.
- This function is supported by DCS Redis 4.0 instances and later. To query the client IP information of Redis 3.0 instances, use the **Client List** command.

# Procedure

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click <sup>(Q)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click a DCS instance to go to the details page.
- **Step 5** Click the **Sessions** tab.
- **Step 6** Information about client connections of the instance is displayed.

#### **NOTE**

- For Proxy Cluster and read/write splitting instances, connections to proxy nodes are displayed. For single-node, master/standby, and Redis Cluster instances, connections to Redis Server nodes are displayed.
- On the page, you can specify a Redis Server or proxy node to query, enter an address, update the query results, and set columns to display.
- If client IP pass-through is disabled for the instance, the value of **addr** is not the actual IP address of the client. Instead, the internal private network IP address **198.19.xxx.xxx** is displayed.
- For details about how to query the actual client IP address, see **Returning the Real IP** Addresses of a Client to DCS (IP Pass-through). After client IP pass-through is enabled, the value of **addr** of new connections is the actual IP address of the client.

# Figure 6-1 Managing sessions

| Sessions Last upo                 | Sessions Last updated: Dec 20, 2023 09:08:32 GMT+08:00 |                    |                   |       |        |      |                         |      |       |        |         |        |        |       |       |        |          |
|-----------------------------------|--|--------------------|-------------------|-------|--------|------|-------------------------|------|-------|--------|---------|--------|--------|-------|-------|--------|----------|
| Kill Selected Kill All • Export • |  |                    |                   |       |        |      |                         |      |       |        |         |        |        |       |       |        |          |
| (group-0)(master)19               | 2.168 🗸 🕻  | Q Select a propert | y or enter a keyw | ord.  |        |      |                         |      |       |        |         |        |        |       |       |        | ) C 🔞    |
| □ ID ÷                            | addr ≑   | name ≑             | cmd 💠             | age ≑ | idle ≑ | db 💠 | flags $\Leftrightarrow$ | fd ≑ | sub ≑ | psub 🔅 | multi 💠 | qbuf 💠 | qbuf ¢ | obl 🔅 | oll ≑ | omem 💠 | events 💠 |
| 215868                            | 192.168  | -                  | set               | 2     | 0      | 0    | N                       | 1930 | 0     | 0      | -1      | 0      | 0      | 0     | 0     | 0      | r        |
| 215869                            | 192.168  | -                  | set               | 2     | 0      | 0    | N                       | 1931 | 0     | 0      | -1      | 0      | 0      | 0     | 0     | 0      | r        |
| 215870                            | 192.168  | -                  | set               | 2     | 0      | 0    | N                       | 1932 | 0     | 0      | -1      | 0      | 0      | 0     | 0     | 0      | r        |
| 215871                            | 192.168  | -                  | set               | 2     | 0      | 0    | N                       | 1933 | 0     | 0      | -1      | 0      | 0      | 0     | 0     | 0      | r        |
| 215872                            | 192.168  |                    | set               | 2     | 0      | 0    | N                       | 1934 | 0     | 0      | -1      | 0      | 0      | 0     | 0     | 0      | r        |
| 215873                            | 192.168  | -                  | set               | 2     | 0      | 0    | N                       | 1935 | 0     | 0      | -1      | 0      | 0      | 0     | 0     | 0      | r        |

Table 6-2 Session fields

| Field     | Description  |
|-----------|--|
| ID        | Unique ID of a session.  |
| addr      | Session address. If IP pass-through is enabled, this address is referred to the client IP address. If not, this address is a private IP address.   |
| name      | Client name, which can be configured using <b>setClientName</b> () in the code. This parameter can be left blank.  |
| cmd       | The last command executed.   |
| age       | Connection duration, in seconds.   |
| idle      | Idle connection duration, in seconds.  |
| db        | The DB identifier in the last executed command, for example, the value of <b>DB0</b> is <b>0</b> .   |
| flags     | Connection flags. <b>M</b> indicates a connection from a master node.<br><b>S</b> indicates a connection from a standby node. For other flags, see https://redis.io/docs/latest/commands/client-list/. |
| fd        | File descriptor.   |
| sub       | Number of channel subscriptions.   |
| psub      | Number of pattern matching subscriptions.  |
| multi     | Number of commands run in transactions or Lua scripts. The value <b>-1</b> indicates that no such command is executed.   |
| qbuf      | Query buffer length (bytes).   |
| qbuf-free | Free space of the query buffer (bytes).  |
| obl       | Output buffer length.  |
| oll       | Output list length.  |
| omem      | Output buffer memory usage (bytes).  |
| events    | File descriptor events (readable, writable). Read: r; Write: w.  |

**Step 7** Select connections to kill and click **Kill Selected** to disconnect the corresponding clients. You can also click **Kill All**.

If a disconnected client can reconnect, it will be automatically reconnected after being disconnected.

Step 8 To export sessions data, click Export. You can export all or selected data.

----End

# 6.4 Modifying Configuration Parameters of a DCS Instance

On the DCS console, you can configure parameters for an instance to achieve optimal DCS performance.

For example, to disable data persistence, set **appendonly** to **no**. For more instance parameters, see **DCS Instance Configuration Parameters**.

#### **NOTE**

After the instance configuration parameters are modified, the modification takes effect immediately without the need to manually restart the instance. For a cluster instance, the modification takes effect on all shards.

### Modifying Configuration Parameters of an Instance

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click <sup>(Q)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** On the **Cache Manager** page, click the name of the DCS instance you want to configure.
- **Step 5** On the instance details page, choose **Instance Configuration** > **Parameters**.
- **Step 6** Click **Modify** in the row containing the desired parameter. To modify multiple parameters at a time, click **Modify** above the parameter list.

Figure 6-2 Modifying parameter(s)

| Parameters Modification History  |  |                     |   |           |
|--|--|---------------------|---|-----------|
| You can modify the parameter values. The new parameter values will take effect | t immediately and apply to all shards of a | a cluster instance. |   |           |
| Modify   |  |                     | Last successful modification: Aug 08, 2024 17:30:22 GMT+08:00 Enter a parameter n | name. Q   |
| Parameter  | Default Value                              | Value Range         | Assigned Value  | Operation |
| active-expire-num (?)  | 20   | 1-1,000             | 20  | Modify    |

#### **Step 7** Modify parameters as required.

The parameters are described in **DCS Instance Configuration Parameters**. In most cases, you can retain default values.

- **Step 8** After you have finished setting the parameters, click **Save**.
- **Step 9** Click **Yes** to confirm the modification.

When the parameter modification task is in the **Successful** state, the parameter is modified.

----End

# **DCS Instance Configuration Parameters**

## 

- For more information about the parameters described in Table 6-3, visit https:// redis.io/topics/memory-optimization.
- Configurable parameters and their values vary depending on the instance type. If a parameter is not displayed in the **Parameters** page on the console, it cannot be modified.

| Parameter             | Description  | Exception<br>Scenario   | Value Range | Default<br>Value |
|-----------------------|--|---|-------------|------------------|
| active-expire-<br>num | Number of<br>randomly<br>checked keys<br>in regular<br>expired key<br>deletions.<br>Enlarging this<br>parameter<br>may increase<br>CPU usage or<br>command<br>latency in a<br>short period<br>of time.<br>Lessening this<br>parameter<br>may increase<br>expired keys<br>in the<br>memory. | This<br>parameter is<br>not available<br>for DCS Redis<br>3.0 and 6.0<br>professional<br>edition<br>instances.<br><b>NOTE</b><br>This<br>parameter<br>was added in<br>September<br>2021. If the<br>parameter<br>value cannot<br>be changed<br>for instances<br>created before<br>September<br>2021, contact<br>customer<br>service. | 1–1000      | 20               |

 Table 6-3 DCS Redis instance configuration parameters

| Parameter | Description  | Exception<br>Scenario | Value Range            | Default<br>Value |
|-----------|--|-----------------------|------------------------|------------------|
| timeout   | The<br>maximum<br>amount of<br>time (in<br>seconds) a<br>connection<br>between a<br>client and the<br>DCS instance<br>can be<br>allowed to<br>remain idle<br>before the<br>connection is<br>terminated.<br>The value <b>0</b><br>indicates that<br>the parameter<br>is disabled.<br>That is, the<br>client is not<br>disconnected<br>when it is<br>idle. |                       | 0–7200<br>Unit: second | 0                |

| Parameter   | Description  | Exception<br>Scenario                                      | Value Range   | Default<br>Value |
|-------------|--|--|---|------------------|
| appendfsync | Controls how<br>often fsync()<br>transfers<br>cached data<br>to the disk.<br>Note that<br>some OSs will<br>perform a<br>complete<br>data transfer<br>but some<br>others only<br>make a "best-<br>effort"<br>attempt. | Single-node<br>instances do<br>not have this<br>parameter. | <ul> <li>no: fsync()<br/>is never<br/>called. The<br/>OS will<br/>flush data<br/>when it is<br/>ready. This<br/>mode<br/>offers the<br/>highest<br/>performan<br/>ce.</li> <li>always:<br/>fsync() is<br/>called after<br/>every write<br/>to the AOF.<br/>This mode<br/>is very<br/>slow, but<br/>also very<br/>safe.</li> <li>everysec:<br/>fsync() is<br/>called once<br/>per second.<br/>This mode<br/>provides a<br/>compromis<br/>e between<br/>safety and<br/>performan<br/>ce.</li> </ul> | no               |

| Parameter  | Description  | Exception<br>Scenario                                      | Value Range  | Default<br>Value |
|------------|--|--|--|------------------|
| appendonly | Indicates<br>whether to<br>log each<br>modification<br>of the<br>instance. By<br>default, data<br>is written to<br>disks<br>asynchronousl<br>y in Redis. If<br>this function<br>is disabled,<br>recently-<br>generated<br>data might be<br>lost in the<br>event of a<br>power failure. | Single-node<br>instances do<br>not have this<br>parameter. | <ul> <li>yes: Logs<br/>are<br/>enabled,<br/>that is,<br/>persistence<br/>is enabled.</li> <li>no: Logs<br/>are<br/>disabled,<br/>that is,<br/>persistence<br/>is disabled.</li> <li>only-<br/>replica:<br/>Enable<br/>persistence<br/>only on<br/>replica<br/>nodes.</li> <li>NOTE<br/>Available in<br/>master/<br/>standby or<br/>cluster DCS<br/>Redis 4.0 and<br/>later basic<br/>edition or<br/>master/<br/>standby<br/>enterprise<br/>(performance)<br/>edition<br/>instances.</li> <li>If this<br/>parameter is<br/>not displayed<br/>on the<br/>console,<br/>contact<br/>customer<br/>service to<br/>enable the<br/>whitelist.</li> </ul> | yes              |

| Parameter   | Description   | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|---|---|--|--|------------------|
| client-output-<br>buffer-limit-<br>slave-soft-<br>seconds | When the<br>client-<br>output-<br>buffer-slave-<br>soft-limit<br>parameter is<br>exceeded for<br>more than<br>the value of<br>this<br>parameter,<br>the server<br>drops the<br>connection.<br>The smaller<br>the value, the<br>easier the<br>disconnection. | Single-node<br>instances do<br>not have this<br>parameter. | 0–60<br>Unit: second                   | 60               |
| client-output-<br>buffer-slave-<br>hard-limit             | Hard limit on<br>the output<br>buffer of<br>replica clients.<br>Once the<br>output buffer<br>exceeds the<br>hard limit, the<br>client is<br>immediately<br>disconnected.<br>The smaller<br>the value, the<br>easier the<br>disconnection.                   | Single-node<br>instances do<br>not have this<br>parameter. | 0–<br>17,179,869,18<br>4<br>Unit: byte | 1,717,986,918    |

| Parameter                                     | Description  | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|---|--|--|--|------------------|
| client-output-<br>buffer-slave-<br>soft-limit | Soft limit on<br>the output<br>buffer of<br>replica clients.<br>Once the<br>output buffer<br>exceeds the<br>soft limit and<br>continuously<br>remains<br>above the<br>limit for the<br>time specified<br>by the client-<br>output-<br>buffer-limit-<br>slave-soft-<br>seconds<br>parameter,<br>the client is<br>disconnected.<br>The smaller<br>the value, the<br>easier the<br>disconnection. | Single-node<br>instances do<br>not have this<br>parameter. | 0–<br>17,179,869,18<br>4<br>Unit: byte | 1,717,986,918    |

| Parameter            | Description  | Exception<br>Scenario | Value Range   | Default<br>Value   |
|----------------------|--|-----------------------|---|--|
| maxmemory-<br>policy | The policy<br>applied when<br>the<br>maxmemory<br>limit is<br>reached. 8<br>values are<br>available. |                       | <ul> <li>volatile-<br/>lru: Evict<br/>keys by<br/>trying to<br/>remove the<br/>less<br/>recently<br/>used (LRU)<br/>keys first,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>lru: Evict<br/>keys by<br/>trying to<br/>remove the<br/>LRU keys<br/>first.</li> <li>volatile-<br/>random:<br/>Evict keys<br/>randomly,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>randomly,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>randomly,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>random:<br/>Evict keys<br/>randomly.</li> <li>volatile-<br/>ttl: Evict<br/>keys with<br/>an expire<br/>set, and try<br/>to evict<br/>keys with a<br/>shorter<br/>time to live<br/>(TTL) first.</li> </ul> | volatile-lru<br>NOTE<br>If the DCS<br>Redis instance<br>is created<br>before July<br>2020 and this<br>parameter has<br>not been<br>modified, the<br>default value<br>is <b>noeviction</b> .<br>If the instance<br>is created<br>after July<br>2020, the<br>default value<br>is <b>volatile-lru</b> . |
|                      |  |                       | delete any<br>keys and  |  |

| Parameter            | Description   | Exception<br>Scenario  | Value Range   | Default<br>Value |
|----------------------|---|--|---|------------------|
|                      |   |  | only return<br>errors<br>when the<br>memory<br>limit was<br>reached.  |                  |
|                      |   |  | <ul> <li>volatile-<br/>lfu: Evict<br/>keys by<br/>trying to<br/>remove the<br/>less<br/>frequently<br/>used (LFU)<br/>keys first,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>lfu: Evict<br/>keys by<br/>trying to<br/>remove the<br/>LFU keys<br/>first.</li> <li>For details<br/>about eviction<br/>policies, see<br/>the Redis<br/>official<br/>website.</li> </ul> |                  |
| lua-time-limit       | Maximum<br>time allowed<br>for executing<br>a Lua script.                         | -  | 100–5,000<br>Unit:<br>millisecond   | 5,000            |
| master-read-<br>only | Sets the<br>instance to be<br>read-only. All<br>write<br>operations<br>will fail. | Proxy Cluster<br>instances do<br>not have this<br>parameter. | <ul><li>yes</li><li>no</li></ul>  | no               |

| Parameter  | Description  | Exception<br>Scenario  | Value Range | Default<br>Value |
|------------|--|--|-------------|------------------|
| maxclients | The maximum number of clients allowed to be concurrently connected to a DCS instance. The larger the value, the more costly the connection to the server, which affects the server performance and increases the command latency. An excessively small value may constrain the server performance. This parameter specifies the maximum number of connections on a single node (single shard). | Read/Write<br>splitting<br>instances do<br>not support<br>this<br>parameter. | 1000-50,000 | 10,000           |

| Parameter              | Description  | Exception<br>Scenario | Value Range                             | Default<br>Value |
|------------------------|--|-----------------------|---|------------------|
|                        | master/<br>standby:<br>Maximum<br>connection<br>s on a<br>single<br>node =<br>Maximum<br>connection<br>s of the<br>instance  |                       |   |                  |
| proto-max-<br>bulk-len | Maximum<br>size of a<br>single<br>element<br>request. Set<br>this<br>parameter to<br>be greater<br>than the<br>customer<br>request<br>length.<br>Otherwise,<br>the request<br>cannot be<br>executed. | -                     | 1,048,576–<br>536,870,912<br>Unit: byte | 536,870,912      |

| Parameter             | Description   | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|-----------------------|---|--|--|------------------|
| repl-backlog-<br>size | The<br>replication<br>backlog size.<br>The backlog is<br>a buffer that<br>accumulates<br>replica data<br>when replicas<br>are<br>disconnected<br>from the<br>master. When<br>a replica<br>reconnects, a<br>partial<br>synchronizatio<br>n is<br>performed to<br>synchronize<br>the data that<br>was missed<br>while replicas<br>were<br>disconnected. | -  | 16,384–<br>1,073,741,824<br>Unit: byte | 1,048,576        |
| repl-backlog-<br>ttl  | The amount<br>of time, in<br>seconds,<br>before the<br>backlog<br>buffer is<br>released,<br>starting from<br>the last a<br>replica was<br>disconnected.<br>The value <b>0</b><br>indicates that<br>the backlog is<br>never<br>released.   | -  | 0–604,800<br>Unit: second              | 3,600            |
| repl-timeout          | Replication<br>timeout.   | Single-node<br>instances do<br>not have this<br>parameter. | 30–3,600<br>Unit: second               | 60               |

| Parameter                    | Description  | Exception<br>Scenario | Value Range | Default<br>Value |
|------------------------------|--|-----------------------|-------------|------------------|
| hash-max-<br>ziplist-entries | The<br>maximum<br>number of<br>hashes that<br>can be<br>encoded<br>using ziplist, a<br>data structure<br>optimized to<br>reduce<br>memory use.   | -                     | 1–10,000    | 512              |
| hash-max-<br>ziplist-value   | The largest<br>value allowed<br>for a hash<br>encoded<br>using ziplist, a<br>special data<br>structure<br>optimized for<br>memory use.   | -                     | 1–10,000    | 64               |
| set-max-<br>intset-entries   | When a set is<br>composed<br>entirely of<br>strings and<br>number of<br>integer<br>elements is<br>less than this<br>parameter<br>value, the set<br>is encoded<br>using intset, a<br>data structure<br>optimized for<br>memory use. | -                     | 1–10,000    | 512              |
| zset-max-<br>ziplist-entries | The<br>maximum<br>number of<br>sorted sets<br>that can be<br>encoded<br>using ziplist, a<br>data structure<br>optimized to<br>reduce<br>memory use.  | -                     | 1–10,000    | 128              |

| Parameter                  | Description  | Exception<br>Scenario | Value Range | Default<br>Value |
|----------------------------|--|-----------------------|-------------|------------------|
| zset-max-<br>ziplist-value | The largest<br>value allowed<br>for a sorted<br>set encoded<br>using ziplist, a<br>special data<br>structure<br>optimized for<br>memory use. | -                     | 1–10,000    | 64               |

| Parameter                         | Description  | Exception<br>Scenario  | Value Range                          | Default<br>Value |
|-----------------------------------|--|--|--------------------------------------|------------------|
| latency-<br>monitor-<br>threshold | The minimum<br>amount of<br>latency that<br>will be logged<br>as latency<br>spikes<br>If this<br>parameter is<br>set to <b>0</b> ,<br>latency<br>monitoring is<br>disabled. If<br>this<br>parameter is<br>set to a value<br>greater than<br>0, all events<br>blocking the<br>server for a<br>time greater<br>than the<br>configured<br>value will be<br>logged.<br>To obtain<br>statistics data,<br>and configure<br>and enable<br>latency<br>monitoring,<br>run the<br><b>LATENCY</b><br>command. | Proxy Cluster<br>instances do<br>not have this<br>parameter. | 0-86,400,000<br>Unit:<br>millisecond | 0                |

| Parameter | Description   | Exception<br>Scenario | Value Range | Default<br>Value |
|-----------|---|-----------------------|-------------|------------------|
|           | NOTE<br>The latency-<br>monitor-<br>threshold<br>parameter is<br>usually used<br>for fault<br>location. After<br>locating faults<br>based on the<br>latency<br>information<br>collected,<br>change the<br>value of<br>latency-<br>monitor-<br>threshold to<br>0 to avoid<br>unnecessary<br>latency. |                       |             |                  |

| Parameter                      | Description  | Exception<br>Scenario  | Value Range   | Default<br>Value |
|--------------------------------|--|--|---|------------------|
| notify-<br>keyspace-<br>events | Controls<br>which<br>keyspace<br>events<br>notifications<br>are enabled<br>for. If this<br>parameter is<br>configured,<br>the Redis<br>Pub/Sub<br>feature will<br>allow clients<br>to receive an<br>event<br>notification<br>when a Redis<br>data set is<br>modified. This<br>parameter is<br>disabled only<br>when it is left<br>blank. | Proxy Cluster<br>instances do<br>not have this<br>parameter. | A<br>combination<br>of different<br>values can be<br>used to<br>enable<br>notifications<br>for multiple<br>event types.<br>Possible<br>values<br>include:<br>K: Keyspace<br>events,<br>published<br>with the<br>keyspace@*<br>prefix<br>E: Keyevent<br>events,<br>published<br>with<br>keyevent@*<br>prefix<br>g: Generic<br>commands<br>(non-type<br>specific) such<br>as DEL,<br>EXPIRE, and<br>RENAME<br>\$: String<br>commands<br>I: List<br>commands<br>S: Set<br>commands<br>s: Set<br>commands<br>x: Expired<br>events<br>(events<br>generated<br>every time a<br>key expires) | Ex               |

| Parameter | Description | Exception<br>Scenario | Value Range   | Default<br>Value |
|-----------|-------------|-----------------------|---|------------------|
|           |             |                       | e: Evicted<br>events<br>(events<br>generated<br>when a key is<br>evicted from<br>maxmemory)   |                  |
|           |             |                       | A: an alias for<br>"g\$lshzxe"  |                  |
|           |             |                       | The<br>parameter<br>value must<br>contain either<br>K or E. A<br>cannot be<br>used together<br>with any of<br>the characters<br>in "g\$lshzxe".<br>For example,<br>the value Kl<br>means that<br>Redis will<br>notify<br>Pub/Sub<br>clients about<br>keyspace<br>events and<br>list<br>commands.<br>The value<br>AKE means<br>Redis will<br>notify<br>Pub/Sub<br>clients about<br>all events. |                  |

| Parameter                             | Description  | Exception<br>Scenario  | Value Range                                  | Default<br>Value |
|---------------------------------------|--|--|--|------------------|
| slowlog-log-<br>slower-than           | Slow queries<br>cover<br>scheduled<br>commands<br>whose<br>execution is<br>delayed.<br><b>slowlog-log-</b><br><b>slower-than</b><br>is the<br>maximum<br>time allowed<br>for command<br>execution. If<br>this threshold<br>is exceeded,<br>Redis will<br>record the<br>query.                                    | -  | 0–1,000,000<br>Unit:<br>microsecond          | 10,000           |
| proxy-<br>slowlog-log-<br>slower-than | Slow queries<br>of a proxy<br>cover<br>scheduled<br>commands<br>whose<br>execution is<br>delayed.<br><b>proxy-</b><br><b>slowlog-log-</b><br><b>slower-than</b><br>is the<br>maximum<br>time allowed<br>for command<br>execution. If<br>this threshold<br>is exceeded,<br>the proxy will<br>record the<br>query. | Currently,<br>only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances in<br>the CN East-<br>Shanghai2<br>and CN<br>South-<br>Guangzhou<br>regions. | 30,000–<br>2,000,000<br>Unit:<br>microsecond | 256,000          |

| Parameter                     | Description  | Exception<br>Scenario  | Value Range | Default<br>Value |
|-------------------------------|--|--|-------------|------------------|
| slowlog-max-<br>len           | The<br>maximum<br>allowed<br>number of<br>slow queries<br>that can be<br>logged. Slow<br>query log<br>consumes<br>memory, but<br>you can<br>reclaim this<br>memory by<br>running the<br><b>SLOWLOG</b><br><b>RESET</b><br>command.               | -  | 0-1000      | 128              |
| proxy-<br>slowlog-max-<br>len | The<br>maximum<br>allowed<br>number of<br>slow queries<br>of a proxy<br>that can be<br>logged. Slow<br>query log<br>consumes<br>memory, but<br>you can<br>reclaim this<br>memory by<br>running the<br><b>SLOWLOG</b><br><b>RESET</b><br>command. | Currently,<br>only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances in<br>the CN East-<br>Shanghai2<br>and CN<br>South-<br>Guangzhou<br>regions. | 0-1000      | 128              |

| Parameter                             | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---------------------------------------|---|--|--|------------------|
| auto-kill-<br>timeout-lua-<br>process | When this<br>parameter is<br>enabled,<br>processes<br>running the<br>lua script are<br>killed when<br>their<br>execution<br>times out.<br>However,<br>scripts with<br>write<br>operations are<br>not killed, but<br>their nodes<br>automatically<br>restart (if<br>persistence<br>has been<br>enabled for<br>the instance)<br>without<br>saving the<br>write<br>operations. | Single-node<br>instances and<br>DCS Redis 3.0<br>instances do<br>not have this<br>parameter. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul> | no               |
| Parameter                               | Description   | Exception<br>Scenario   | Value Range  | Default<br>Value |
|---|---|---|--|------------------|
| audit-log-<br>customer-<br>command-list | Commands to<br>record in<br>audit logs<br>(only write<br>commands<br>are recorded<br>by default.)<br>This<br>parameter is<br>valid only<br>when the<br>audit log<br>function is<br>enabled. | Viewing<br>Audit Logs of<br>a DCS Redis<br>Instance is<br>available only<br>in certain<br>regions. This<br>parameter is<br>displayed only<br>for Proxy<br>Cluster<br>instances<br>when the<br>audit log<br>feature is<br>supported. | A maximum<br>of 10<br>commands<br>are allowed.<br>For each<br>command,<br>use up to 10<br>characters<br>including<br>letters,<br>periods (.),<br>hyphens (-),<br>and<br>underscores<br>(_), and start<br>and end with<br>a letter.<br>Separate<br>multiple<br>commands<br>with spaces,<br>and end your<br>input with a<br>space. |                  |

| Parameter               | Description   | Exception<br>Scenario   | Value Range  | Default<br>Value |
|-------------------------|---|---|--|------------------|
| backend-<br>master-only | Read/Write<br>splitting is<br>disabled by<br>default for<br>Proxy Cluster<br>instances. In<br>this case, read<br>and write<br>requests are<br>allocated to<br>the master<br>node of a<br>Proxy Cluster<br>instance.<br>For Proxy                          | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter. | <ul> <li>yes:<br/>disables<br/>read/write<br/>splitting.</li> <li>no:<br/>enables<br/>read/write<br/>splitting.</li> </ul> | yes              |
|                         | Cluster<br>instances with<br>read/write<br>splitting<br>enabled, read<br>requests are<br>allocated<br>evenly to<br>each master<br>or standby<br>node of a<br>Proxy Cluster<br>instance while<br>write requests<br>are allocated<br>to the master<br>node. |   |  |                  |

| Parameter                             | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---------------------------------------|---|--|--|------------------|
| read-only-<br>slave-when-<br>wr-split | By default,<br>Proxy Cluster<br>instances<br>perform read<br>operations<br>only on<br>standby<br>nodes. To<br>perform read<br>operations on<br>both master<br>and standby<br>nodes,<br>configure this<br>parameter. | Only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances<br>have this<br>parameter<br>and it is only<br>applicable<br>when read/<br>write splitting<br>is enabled.<br>Read/Write<br>splitting is<br>enabled by<br>default for<br>read/write<br>splitting<br>instances. For<br>Proxy Cluster<br>instances, it<br>can be<br>enabled by<br>setting<br>parameter<br><b>backend-</b><br><b>master-only</b><br>to <b>no</b> . | <b>yes</b> : Read<br>only on<br>standby<br>nodes.<br><b>no</b> : Read on<br>both master<br>and standby<br>nodes. | yes              |

| Parameter                                   | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---|---|--|--|------------------|
| dispatch-<br>pubsub-to-<br>fixed-shard      | This<br>parameter<br>specifies<br>whether<br>pub/sub<br>channels are<br>on the shard<br>of slot 0.<br>When this<br>parameter is<br>enabled, the<br>pub/sub<br>processing<br>logic is<br>consistent<br>with that of<br>single-node<br>instances. You<br>are advised to<br>enable this<br>parameter if<br>you do not<br>depend<br>heavily on<br>pub/sub. If<br>you depend<br>heavily on<br>pub/sub, use<br>the default<br>configuration<br>to allocate<br>subscriptions<br>to all shards. | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter.            | <ul> <li>yes: Enable this parameter to allocate subscriptio n channels to the shard of slot 0.</li> <li>no: Disable this parameter to allocate channels to the shard of each channel-hashed slot.</li> </ul> | no               |
| readonly-lua-<br>route-to-<br>slave-enabled | If enabled,<br>read-only Lua<br>scripts of<br>read-only<br>users are<br>executed and<br>routed to the<br>standby node.  | Only read/<br>write splitting<br>instances<br>support this<br>parameter. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul>   | no               |
| cluster-<br>sentinel-<br>enabled            | To support<br>Sentinels for<br>the instance.  | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter.            | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul>   | no               |

| Parameter                 | Description   | Exception<br>Scenario   | Value Range  | Default<br>Value |
|---------------------------|---|---|--|------------------|
| scan-support-<br>wr-split | The <b>SCAN</b><br>command is<br>executed on<br>the master<br>node when<br>this<br>parameter is<br>disabled, or is<br>executed on<br>the standby<br>node<br>otherwise.<br>Enabling this<br>parameter<br>relieves SCAN<br>commands on<br>the master<br>node. But<br>newly written<br>data in the<br>master node<br>may not be<br>synchronized<br>to replicas in<br>time. | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter.<br>Proxy Cluster<br>instances<br>created<br>earlier may<br>not support<br>this<br>parameter. In<br>this case,<br>contact<br>customer<br>service to<br>upgrade<br>instances. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul> | no               |

# Table 6-4 DCS Memcached instance configuration parameters

| Parameter  | Description   | Value Range            | Default<br>Value |
|------------|---|------------------------|------------------|
| timeout    | The maximum amount of<br>time (in seconds) a<br>connection between a<br>client and the DCS instance<br>can be allowed to remain<br>idle before the connection<br>is terminated. A setting of<br><b>0</b> means that this function<br>is disabled. | 0–7200<br>Unit: second | 0                |
| maxclients | The maximum number of clients allowed to be concurrently connected to a DCS instance.   | 1000–10,000            | 10,000           |

| Parameter                       | Description  | Value Range  | Default<br>Value |
|---------------------------------|--|--|------------------|
| maxmemor<br>y-policy            | The policy applied when<br>the maxmemory limit is<br>reached.  | <ul> <li>volatile-lru: Evict<br/>keys by trying to<br/>remove the less<br/>recently used (LRU)<br/>keys first, but only<br/>among keys that<br/>have an expire set.</li> <li>allkeys-lru: Evict<br/>keys by trying to<br/>remove the LRU keys<br/>first.</li> <li>volatile-random:<br/>Evict keys randomly,</li> </ul> | noevictio<br>n   |
|                                 |  | that have an expire set.   |                  |
|                                 |  | • allkeys-random:<br>Evict keys randomly.  |                  |
|                                 |  | <ul> <li>volatile-ttl: Evict<br/>keys with an expire<br/>set, and try to evict<br/>keys with a shorter<br/>time to live (TTL)<br/>first.</li> </ul>  |                  |
|                                 |  | • <b>noeviction</b> : Do not delete any keys and only return errors when the memory limit was reached.   |                  |
| reserved-<br>memory-<br>percent | Percentage of the<br>maximum available<br>memory reserved for<br>background processes, such<br>as data persistence and<br>replication. | 0–80   | 30               |

# 6.5 Configuring DCS Instance Parameter Templates

# 6.5.1 Viewing a Parameter Template of a DCS Instance

System default parameter templates vary by Redis version and instance type. A system default parameter template contains default instance parameter configurations. Parameter templates can be customized for parameter configurations, and can be selected in instance creation.

This section describes how to view instance parameter templates on the DCS console.

# Procedure

- **Step 1** Log in to the DCS console.
- **Step 2** Click O in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Parameter Templates**.
- **Step 4** Choose the **Default Templates** or **Custom Templates** tab.
- **Step 5** View parameter templates.

Currently, you can enter a keyword in the search box to search for a parameter template by template name.

**Step 6** Click a parameter template. The parameters contained in the template are displayed. For details about the parameters, see **Table 6-5**.

| Parameter             | Description  | Exception<br>Scenario   | Value Range | Default<br>Value |
|-----------------------|--|---|-------------|------------------|
| active-expire-<br>num | Number of<br>randomly<br>checked keys<br>in regular<br>expired key<br>deletions.<br>Enlarging this<br>parameter<br>may increase<br>CPU usage or<br>command<br>latency in a<br>short period<br>of time.<br>Lessening this<br>parameter<br>may increase<br>expired keys<br>in the<br>memory. | This<br>parameter is<br>not available<br>for DCS Redis<br>3.0 and 6.0<br>professional<br>edition<br>instances.<br><b>NOTE</b><br>This<br>parameter<br>was added in<br>September<br>2021. If the<br>parameter<br>value cannot<br>be changed<br>for instances<br>created before<br>September<br>2021, contact<br>customer<br>service. | 1–1000      | 20               |

**Table 6-5** DCS Redis instance configuration parameters

| Parameter | Description  | Exception<br>Scenario | Value Range            | Default<br>Value |
|-----------|--|-----------------------|------------------------|------------------|
| timeout   | The<br>maximum<br>amount of<br>time (in<br>seconds) a<br>connection<br>between a<br>client and the<br>DCS instance<br>can be<br>allowed to<br>remain idle<br>before the<br>connection is<br>terminated.<br>The value <b>0</b><br>indicates that<br>the parameter<br>is disabled.<br>That is, the<br>client is not<br>disconnected<br>when it is<br>idle. |                       | 0–7200<br>Unit: second | 0                |

| Parameter   | Description  | Exception<br>Scenario                                      | Value Range   | Default<br>Value |
|-------------|--|--|---|------------------|
| appendfsync | Controls how<br>often fsync()<br>transfers<br>cached data<br>to the disk.<br>Note that<br>some OSs will<br>perform a<br>complete<br>data transfer<br>but some<br>others only<br>make a "best-<br>effort"<br>attempt. | Single-node<br>instances do<br>not have this<br>parameter. | <ul> <li>no: fsync()<br/>is never<br/>called. The<br/>OS will<br/>flush data<br/>when it is<br/>ready. This<br/>mode<br/>offers the<br/>highest<br/>performan<br/>ce.</li> <li>always:<br/>fsync() is<br/>called after<br/>every write<br/>to the AOF.<br/>This mode<br/>is very<br/>slow, but<br/>also very<br/>safe.</li> <li>everysec:<br/>fsync() is<br/>called once<br/>per second.<br/>This mode<br/>provides a<br/>compromis<br/>e between<br/>safety and<br/>performan<br/>ce.</li> </ul> | no               |

| Parameter  | Description  | Exception<br>Scenario                                      | Value Range  | Default<br>Value |
|------------|--|--|--|------------------|
| appendonly | Indicates<br>whether to<br>log each<br>modification<br>of the<br>instance. By<br>default, data<br>is written to<br>disks<br>asynchronousl<br>y in Redis. If<br>this function<br>is disabled,<br>recently-<br>generated<br>data might be<br>lost in the<br>event of a<br>power failure. | Single-node<br>instances do<br>not have this<br>parameter. | <ul> <li>yes: Logs<br/>are<br/>enabled,<br/>that is,<br/>persistence<br/>is enabled.</li> <li>no: Logs<br/>are<br/>disabled,<br/>that is,<br/>persistence<br/>is disabled.</li> <li>only-<br/>replica:<br/>Enable<br/>persistence<br/>only on<br/>replica<br/>nodes.</li> <li>NOTE<br/>Available in<br/>master/<br/>standby or<br/>cluster DCS<br/>Redis 4.0 and<br/>later basic<br/>edition or<br/>master/<br/>standby<br/>enterprise<br/>(performance)<br/>edition<br/>instances.</li> <li>If this<br/>parameter is<br/>not displayed<br/>on the<br/>console,<br/>contact<br/>customer<br/>service to<br/>enable the<br/>whitelist.</li> </ul> | yes              |

| Parameter   | Description   | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|---|---|--|--|------------------|
| client-output-<br>buffer-limit-<br>slave-soft-<br>seconds | When the<br>client-<br>output-<br>buffer-slave-<br>soft-limit<br>parameter is<br>exceeded for<br>more than<br>the value of<br>this<br>parameter,<br>the server<br>drops the<br>connection.<br>The smaller<br>the value, the<br>easier the<br>disconnection. | Single-node<br>instances do<br>not have this<br>parameter. | 0–60<br>Unit: second                   | 60               |
| client-output-<br>buffer-slave-<br>hard-limit             | Hard limit on<br>the output<br>buffer of<br>replica clients.<br>Once the<br>output buffer<br>exceeds the<br>hard limit, the<br>client is<br>immediately<br>disconnected.<br>The smaller<br>the value, the<br>easier the<br>disconnection.                   | Single-node<br>instances do<br>not have this<br>parameter. | 0–<br>17,179,869,18<br>4<br>Unit: byte | 1,717,986,918    |

| Parameter                                     | Description  | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|---|--|--|--|------------------|
| client-output-<br>buffer-slave-<br>soft-limit | Soft limit on<br>the output<br>buffer of<br>replica clients.<br>Once the<br>output buffer<br>exceeds the<br>soft limit and<br>continuously<br>remains<br>above the<br>limit for the<br>time specified<br>by the client-<br>output-<br>buffer-limit-<br>slave-soft-<br>seconds<br>parameter,<br>the client is<br>disconnected.<br>The smaller<br>the value, the<br>easier the<br>disconnection. | Single-node<br>instances do<br>not have this<br>parameter. | 0–<br>17,179,869,18<br>4<br>Unit: byte | 1,717,986,918    |

| Parameter            | Description  | Exception<br>Scenario | Value Range   | Default<br>Value   |
|----------------------|--|-----------------------|---|--|
| maxmemory-<br>policy | The policy<br>applied when<br>the<br>maxmemory<br>limit is<br>reached. 8<br>values are<br>available. |                       | <ul> <li>volatile-<br/>lru: Evict<br/>keys by<br/>trying to<br/>remove the<br/>less<br/>recently<br/>used (LRU)<br/>keys first,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>lru: Evict<br/>keys by<br/>trying to<br/>remove the<br/>LRU keys<br/>first.</li> <li>volatile-<br/>random:<br/>Evict keys<br/>randomly,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>randomly,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>randomly,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>random:<br/>Evict keys<br/>randomly.</li> <li>volatile-<br/>ttl: Evict<br/>keys with<br/>an expire<br/>set, and try<br/>to evict<br/>keys with a<br/>shorter<br/>time to live<br/>(TTL) first.</li> </ul> | volatile-lru<br>NOTE<br>If the DCS<br>Redis instance<br>is created<br>before July<br>2020 and this<br>parameter has<br>not been<br>modified, the<br>default value<br>is <b>noeviction</b> .<br>If the instance<br>is created<br>after July<br>2020, the<br>default value<br>is <b>volatile-lru</b> . |
|                      |  |                       | delete any<br>keys and  |  |

| Parameter            | Description   | Exception<br>Scenario  | Value Range   | Default<br>Value |
|----------------------|---|--|---|------------------|
|                      |   |  | only return<br>errors<br>when the<br>memory<br>limit was<br>reached.  |                  |
|                      |   |  | <ul> <li>volatile-<br/>lfu: Evict<br/>keys by<br/>trying to<br/>remove the<br/>less<br/>frequently<br/>used (LFU)<br/>keys first,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>lfu: Evict<br/>keys by<br/>trying to<br/>remove the<br/>LFU keys<br/>first.</li> <li>For details<br/>about eviction<br/>policies, see<br/>the Redis<br/>official<br/>website.</li> </ul> |                  |
| lua-time-limit       | Maximum<br>time allowed<br>for executing<br>a Lua script.                         | -  | 100–5,000<br>Unit:<br>millisecond   | 5,000            |
| master-read-<br>only | Sets the<br>instance to be<br>read-only. All<br>write<br>operations<br>will fail. | Proxy Cluster<br>instances do<br>not have this<br>parameter. | <ul><li>yes</li><li>no</li></ul>  | no               |

| Parameter  | Description  | Exception<br>Scenario  | Value Range | Default<br>Value |
|------------|--|--|-------------|------------------|
| maxclients | The maximum number of clients allowed to be concurrently connected to a DCS instance. The larger the value, the more costly the connection to the server, which affects the server performance and increases the command latency. An excessively small value may constrain the server performance. This parameter specifies the maximum number of connections on a single node (single shard). | Read/Write<br>splitting<br>instances do<br>not support<br>this<br>parameter. | 1000-50,000 | 10,000           |
|            | node and   |  |             |                  |

| Parameter              | Description  | Exception<br>Scenario | Value Range                             | Default<br>Value |
|------------------------|--|-----------------------|---|------------------|
|                        | master/<br>standby:<br>Maximum<br>connection<br>s on a<br>single<br>node =<br>Maximum<br>connection<br>s of the<br>instance  |                       |   |                  |
| proto-max-<br>bulk-len | Maximum<br>size of a<br>single<br>element<br>request. Set<br>this<br>parameter to<br>be greater<br>than the<br>customer<br>request<br>length.<br>Otherwise,<br>the request<br>cannot be<br>executed. | -                     | 1,048,576–<br>536,870,912<br>Unit: byte | 536,870,912      |

| Parameter             | Description   | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|-----------------------|---|--|--|------------------|
| repl-backlog-<br>size | The<br>replication<br>backlog size.<br>The backlog is<br>a buffer that<br>accumulates<br>replica data<br>when replicas<br>are<br>disconnected<br>from the<br>master. When<br>a replica<br>reconnects, a<br>partial<br>synchronizatio<br>n is<br>performed to<br>synchronize<br>the data that<br>was missed<br>while replicas<br>were<br>disconnected. | -  | 16,384–<br>1,073,741,824<br>Unit: byte | 1,048,576        |
| repl-backlog-<br>ttl  | The amount<br>of time, in<br>seconds,<br>before the<br>backlog<br>buffer is<br>released,<br>starting from<br>the last a<br>replica was<br>disconnected.<br>The value <b>0</b><br>indicates that<br>the backlog is<br>never<br>released.   | -  | 0–604,800<br>Unit: second              | 3,600            |
| repl-timeout          | Replication<br>timeout.   | Single-node<br>instances do<br>not have this<br>parameter. | 30–3,600<br>Unit: second               | 60               |

| Parameter                    | Description  | Exception<br>Scenario | Value Range | Default<br>Value |
|------------------------------|--|-----------------------|-------------|------------------|
| hash-max-<br>ziplist-entries | The<br>maximum<br>number of<br>hashes that<br>can be<br>encoded<br>using ziplist, a<br>data structure<br>optimized to<br>reduce<br>memory use.   | -                     | 1–10,000    | 512              |
| hash-max-<br>ziplist-value   | The largest<br>value allowed<br>for a hash<br>encoded<br>using ziplist, a<br>special data<br>structure<br>optimized for<br>memory use.   | -                     | 1–10,000    | 64               |
| set-max-<br>intset-entries   | When a set is<br>composed<br>entirely of<br>strings and<br>number of<br>integer<br>elements is<br>less than this<br>parameter<br>value, the set<br>is encoded<br>using intset, a<br>data structure<br>optimized for<br>memory use. | -                     | 1–10,000    | 512              |
| zset-max-<br>ziplist-entries | The<br>maximum<br>number of<br>sorted sets<br>that can be<br>encoded<br>using ziplist, a<br>data structure<br>optimized to<br>reduce<br>memory use.  | -                     | 1–10,000    | 128              |

| Parameter                  | Description  | Exception<br>Scenario | Value Range | Default<br>Value |
|----------------------------|--|-----------------------|-------------|------------------|
| zset-max-<br>ziplist-value | The largest<br>value allowed<br>for a sorted<br>set encoded<br>using ziplist, a<br>special data<br>structure<br>optimized for<br>memory use. | -                     | 1–10,000    | 64               |

| Parameter                         | Description  | Exception<br>Scenario  | Value Range                          | Default<br>Value |
|-----------------------------------|--|--|--------------------------------------|------------------|
| latency-<br>monitor-<br>threshold | The minimum<br>amount of<br>latency that<br>will be logged<br>as latency<br>spikes   | Proxy Cluster<br>instances do<br>not have this<br>parameter. | 0–86,400,000<br>Unit:<br>millisecond | 0                |
|                                   | If this<br>parameter is<br>set to <b>0</b> ,<br>latency<br>monitoring is<br>disabled. If<br>this<br>parameter is<br>set to a value<br>greater than<br>0, all events<br>blocking the<br>server for a<br>time greater<br>than the<br>configured<br>value will be<br>logged.<br>To obtain<br>statistics data,<br>and configure<br>and enable<br>latency<br>monitoring,<br>run the<br><b>LATENCY</b><br>command. |  |                                      |                  |

| Parameter | Description   | Exception<br>Scenario | Value Range | Default<br>Value |
|-----------|---|-----------------------|-------------|------------------|
|           | NOTE<br>The latency-<br>monitor-<br>threshold<br>parameter is<br>usually used<br>for fault<br>location. After<br>locating faults<br>based on the<br>latency<br>information<br>collected,<br>change the<br>value of<br>latency-<br>monitor-<br>threshold to<br>0 to avoid<br>unnecessary<br>latency. |                       |             |                  |

| Parameter                      | Description  | Exception<br>Scenario  | Value Range   | Default<br>Value |
|--------------------------------|--|--|---|------------------|
| notify-<br>keyspace-<br>events | Controls<br>which<br>keyspace<br>events<br>notifications<br>are enabled<br>for. If this<br>parameter is<br>configured,<br>the Redis<br>Pub/Sub<br>feature will<br>allow clients<br>to receive an<br>event<br>notification<br>when a Redis<br>data set is<br>modified. This<br>parameter is<br>disabled only<br>when it is left<br>blank. | Proxy Cluster<br>instances do<br>not have this<br>parameter. | A<br>combination<br>of different<br>values can be<br>used to<br>enable<br>notifications<br>for multiple<br>event types.<br>Possible<br>values<br>include:<br>K: Keyspace<br>events,<br>published<br>with the<br>keyspace@*<br>prefix<br>E: Keyevent<br>events,<br>published<br>with<br>keyevent@*<br>prefix<br>g: Generic<br>commands<br>(non-type<br>specific) such<br>as DEL,<br>EXPIRE, and<br>RENAME<br>\$: String<br>commands<br>I: List<br>commands<br>S: Set<br>commands<br>s: Set<br>commands<br>x: Expired<br>events<br>(events<br>generated<br>every time a<br>key expires) | Ex               |

| Parameter | Description | Exception<br>Scenario | Value Range   | Default<br>Value |
|-----------|-------------|-----------------------|---|------------------|
|           |             |                       | e: Evicted<br>events<br>(events<br>generated<br>when a key is<br>evicted from<br>maxmemory)   |                  |
|           |             |                       | A: an alias for<br>"g\$lshzxe"  |                  |
|           |             |                       | The<br>parameter<br>value must<br>contain either<br>K or E. A<br>cannot be<br>used together<br>with any of<br>the characters<br>in "g\$lshzxe".<br>For example,<br>the value Kl<br>means that<br>Redis will<br>notify<br>Pub/Sub<br>clients about<br>keyspace<br>events and<br>list<br>commands.<br>The value<br>AKE means<br>Redis will<br>notify<br>Pub/Sub<br>clients about<br>all events. |                  |

| Parameter                             | Description  | Exception<br>Scenario  | Value Range                                  | Default<br>Value |
|---------------------------------------|--|--|--|------------------|
| slowlog-log-<br>slower-than           | Slow queries<br>cover<br>scheduled<br>commands<br>whose<br>execution is<br>delayed.<br><b>slowlog-log-</b><br><b>slower-than</b><br>is the<br>maximum<br>time allowed<br>for command<br>execution. If<br>this threshold<br>is exceeded,<br>Redis will<br>record the<br>query.                                    | -  | 0–1,000,000<br>Unit:<br>microsecond          | 10,000           |
| proxy-<br>slowlog-log-<br>slower-than | Slow queries<br>of a proxy<br>cover<br>scheduled<br>commands<br>whose<br>execution is<br>delayed.<br><b>proxy-</b><br><b>slowlog-log-</b><br><b>slower-than</b><br>is the<br>maximum<br>time allowed<br>for command<br>execution. If<br>this threshold<br>is exceeded,<br>the proxy will<br>record the<br>query. | Currently,<br>only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances in<br>the CN East-<br>Shanghai2<br>and CN<br>South-<br>Guangzhou<br>regions. | 30,000–<br>2,000,000<br>Unit:<br>microsecond | 256,000          |

| Parameter                     | Description  | Exception<br>Scenario  | Value Range | Default<br>Value |
|-------------------------------|--|--|-------------|------------------|
| slowlog-max-<br>len           | The<br>maximum<br>allowed<br>number of<br>slow queries<br>that can be<br>logged. Slow<br>query log<br>consumes<br>memory, but<br>you can<br>reclaim this<br>memory by<br>running the<br><b>SLOWLOG</b><br><b>RESET</b><br>command.               | -  | 0-1000      | 128              |
| proxy-<br>slowlog-max-<br>len | The<br>maximum<br>allowed<br>number of<br>slow queries<br>of a proxy<br>that can be<br>logged. Slow<br>query log<br>consumes<br>memory, but<br>you can<br>reclaim this<br>memory by<br>running the<br><b>SLOWLOG</b><br><b>RESET</b><br>command. | Currently,<br>only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances in<br>the CN East-<br>Shanghai2<br>and CN<br>South-<br>Guangzhou<br>regions. | 0-1000      | 128              |

| Parameter                             | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---------------------------------------|---|--|--|------------------|
| auto-kill-<br>timeout-lua-<br>process | When this<br>parameter is<br>enabled,<br>processes<br>running the<br>lua script are<br>killed when<br>their<br>execution<br>times out.<br>However,<br>scripts with<br>write<br>operations are<br>not killed, but<br>their nodes<br>automatically<br>restart (if<br>persistence<br>has been<br>enabled for<br>the instance)<br>without<br>saving the<br>write<br>operations. | Single-node<br>instances and<br>DCS Redis 3.0<br>instances do<br>not have this<br>parameter. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul> | no               |

| Parameter                               | Description   | Exception<br>Scenario   | Value Range  | Default<br>Value |
|---|---|---|--|------------------|
| audit-log-<br>customer-<br>command-list | Commands to<br>record in<br>audit logs<br>(only write<br>commands<br>are recorded<br>by default.)<br>This<br>parameter is<br>valid only<br>when the<br>audit log<br>function is<br>enabled. | Viewing<br>Audit Logs of<br>a DCS Redis<br>Instance is<br>available only<br>in certain<br>regions. This<br>parameter is<br>displayed only<br>for Proxy<br>Cluster<br>instances<br>when the<br>audit log<br>feature is<br>supported. | A maximum<br>of 10<br>commands<br>are allowed.<br>For each<br>command,<br>use up to 10<br>characters<br>including<br>letters,<br>periods (.),<br>hyphens (-),<br>and<br>underscores<br>(_), and start<br>and end with<br>a letter.<br>Separate<br>multiple<br>commands<br>with spaces,<br>and end your<br>input with a<br>space. |                  |

| Parameter               | Description  | Exception<br>Scenario   | Value Range  | Default<br>Value |
|-------------------------|--|---|--|------------------|
| backend-<br>master-only | Read/Write<br>splitting is<br>disabled by<br>default for<br>Proxy Cluster<br>instances. In<br>this case, read<br>and write<br>requests are<br>allocated to<br>the master<br>node of a<br>Proxy Cluster<br>instance.<br>For Proxy<br>Cluster<br>instances with<br>read/write<br>splitting | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter. | <ul> <li>yes:<br/>disables<br/>read/write<br/>splitting.</li> <li>no:<br/>enables<br/>read/write<br/>splitting.</li> </ul> | yes              |
|                         | enabled, read<br>requests are<br>allocated<br>evenly to<br>each master<br>or standby<br>node of a<br>Proxy Cluster<br>instance while<br>write requests<br>are allocated<br>to the master<br>node   |   |  |                  |

| Parameter                             | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---------------------------------------|---|--|--|------------------|
| read-only-<br>slave-when-<br>wr-split | By default,<br>Proxy Cluster<br>instances<br>perform read<br>operations<br>only on<br>standby<br>nodes. To<br>perform read<br>operations on<br>both master<br>and standby<br>nodes,<br>configure this<br>parameter. | Only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances<br>have this<br>parameter<br>and it is only<br>applicable<br>when read/<br>write splitting<br>is enabled.<br>Read/Write<br>splitting is<br>enabled by<br>default for<br>read/write<br>splitting<br>instances. For<br>Proxy Cluster<br>instances, it<br>can be<br>enabled by<br>setting<br>parameter<br><b>backend-</b><br><b>master-only</b><br>to <b>no</b> . | <b>yes</b> : Read<br>only on<br>standby<br>nodes.<br><b>no</b> : Read on<br>both master<br>and standby<br>nodes. | yes              |

| Parameter                                   | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---|---|--|--|------------------|
| dispatch-<br>pubsub-to-<br>fixed-shard      | This<br>parameter<br>specifies<br>whether<br>pub/sub<br>channels are<br>on the shard<br>of slot 0.<br>When this<br>parameter is<br>enabled, the<br>pub/sub<br>processing<br>logic is<br>consistent<br>with that of<br>single-node<br>instances. You<br>are advised to<br>enable this<br>parameter if<br>you do not<br>depend<br>heavily on<br>pub/sub. If<br>you depend<br>heavily on<br>pub/sub, use<br>the default<br>configuration<br>to allocate<br>subscriptions<br>to all shards. | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter.            | <ul> <li>yes: Enable this parameter to allocate subscriptio n channels to the shard of slot 0.</li> <li>no: Disable this parameter to allocate channels to the shard of each channel-hashed slot.</li> </ul> | no               |
| readonly-lua-<br>route-to-<br>slave-enabled | If enabled,<br>read-only Lua<br>scripts of<br>read-only<br>users are<br>executed and<br>routed to the<br>standby node.  | Only read/<br>write splitting<br>instances<br>support this<br>parameter. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul>   | no               |
| cluster-<br>sentinel-<br>enabled            | To support<br>Sentinels for<br>the instance.  | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter.            | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul>   | no               |

| Parameter                 | Description   | Exception<br>Scenario   | Value Range  | Default<br>Value |
|---------------------------|---|---|--|------------------|
| scan-support-<br>wr-split | The <b>SCAN</b><br>command is<br>executed on<br>the master<br>node when<br>this<br>parameter is<br>disabled, or is<br>executed on<br>the standby<br>node<br>otherwise.<br>Enabling this<br>parameter<br>relieves SCAN<br>commands on<br>the master<br>node. But<br>newly written<br>data in the<br>master node<br>may not be<br>synchronized<br>to replicas in<br>time. | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter.<br>Proxy Cluster<br>instances<br>created<br>earlier may<br>not support<br>this<br>parameter. In<br>this case,<br>contact<br>customer<br>service to<br>upgrade<br>instances. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul> | no               |

## D NOTE

- 1. The default values and value ranges of the **maxclients**, **reserved-memory-percent**, **client-output-buffer-slave-soft-limit**, and **client-output-buffer-slave-hard-limit** parameters are related to the instance specifications. Therefore, these parameters are not displayed in the parameter template.
- 2. For more information about the parameters described in Table 6-5, visit https:// redis.io/topics/memory-optimization.

### ----End

# 6.5.2 Creating a Custom Parameter Template for a DCS Instance

System default parameter templates vary by Redis version and instance type. A system default parameter template contains default instance parameter configurations. Parameter templates can be customized for parameter configurations, and can be selected in instance creation.

This section describes how to create and modify a custom parameter template on the DCS console.

## Procedure

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

- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Parameter Templates**.
- **Step 4** Click the **Default Templates** or **Custom Templates** tab to create a template based on a default template or an existing custom template.
  - If you select **Default Templates**, click **Customize** in the **Operation** column of the row containing the desired cache engine version.
  - If you select **Custom Templates**, click **Copy** in the **Operation** column in the row containing the desired custom template.

### Step 5 Specify Template Name and Description.

### 

The template name can contain 4 to 64 characters and must start with a letter or digit. Only letters, digits, hyphens (-), underscores (\_), and periods (.) are allowed. The description can be empty.

### Step 6 Select Modifiable parameters.

Currently, you can enter a keyword in the search box to search for a parameter by parameter name.

**Step 7** In the row that contains the parameter to be modified, enter a value in the **Assigned Value** column.

 Table 6-6 describes the parameters. In most cases, default values are retained.

| Parameter             | Description  | Exception<br>Scenario   | Value Range            | Default<br>Value |
|-----------------------|--|---|------------------------|------------------|
| active-expire-<br>num | Number of<br>randomly<br>checked keys<br>in regular<br>expired key<br>deletions.<br>Enlarging this<br>parameter<br>may increase<br>CPU usage or<br>command<br>latency in a<br>short period<br>of time.<br>Lessening this<br>parameter<br>may increase<br>expired keys<br>in the<br>memory.   | This<br>parameter is<br>not available<br>for DCS Redis<br>3.0 and 6.0<br>professional<br>edition<br>instances.<br><b>NOTE</b><br>This<br>parameter<br>was added in<br>September<br>2021. If the<br>parameter<br>value cannot<br>be changed<br>for instances<br>created before<br>September<br>2021, contact<br>customer<br>service. | 1-1000                 | 20               |
| timeout               | The<br>maximum<br>amount of<br>time (in<br>seconds) a<br>connection<br>between a<br>client and the<br>DCS instance<br>can be<br>allowed to<br>remain idle<br>before the<br>connection is<br>terminated.<br>The value <b>0</b><br>indicates that<br>the parameter<br>is disabled.<br>That is, the<br>client is not<br>disconnected<br>when it is<br>idle. |   | 0–7200<br>Unit: second | 0                |

 Table 6-6 DCS Redis instance configuration parameters

| Parameter   | Description  | Exception<br>Scenario                                      | Value Range   | Default<br>Value |
|-------------|--|--|---|------------------|
| appendfsync | Controls how<br>often fsync()<br>transfers<br>cached data<br>to the disk.<br>Note that<br>some OSs will<br>perform a<br>complete<br>data transfer<br>but some<br>others only<br>make a "best-<br>effort"<br>attempt. | Single-node<br>instances do<br>not have this<br>parameter. | <ul> <li>no: fsync()<br/>is never<br/>called. The<br/>OS will<br/>flush data<br/>when it is<br/>ready. This<br/>mode<br/>offers the<br/>highest<br/>performan<br/>ce.</li> <li>always:<br/>fsync() is<br/>called after<br/>every write<br/>to the AOF.<br/>This mode<br/>is very<br/>slow, but<br/>also very<br/>safe.</li> <li>everysec:<br/>fsync() is<br/>called once<br/>per second.<br/>This mode<br/>provides a<br/>compromis<br/>e between<br/>safety and<br/>performan<br/>ce.</li> </ul> | no               |

| Parameter  | Description  | Exception<br>Scenario                                      | Value Range  | Default<br>Value |
|------------|--|--|--|------------------|
| appendonly | Indicates<br>whether to<br>log each<br>modification<br>of the<br>instance. By<br>default, data<br>is written to<br>disks<br>asynchronousl<br>y in Redis. If<br>this function<br>is disabled,<br>recently-<br>generated<br>data might be<br>lost in the<br>event of a<br>power failure. | Single-node<br>instances do<br>not have this<br>parameter. | <ul> <li>yes: Logs<br/>are<br/>enabled,<br/>that is,<br/>persistence<br/>is enabled.</li> <li>no: Logs<br/>are<br/>disabled,<br/>that is,<br/>persistence<br/>is disabled.</li> <li>only-<br/>replica:<br/>Enable<br/>persistence<br/>only on<br/>replica<br/>nodes.</li> <li>NOTE<br/>Available in<br/>master/<br/>standby or<br/>cluster DCS<br/>Redis 4.0 and<br/>later basic<br/>edition or<br/>master/<br/>standby<br/>enterprise<br/>(performance)<br/>edition<br/>instances.</li> <li>If this<br/>parameter is<br/>not displayed<br/>on the<br/>console,<br/>contact<br/>customer<br/>service to<br/>enable the<br/>whitelist.</li> </ul> | yes              |

| Parameter   | Description   | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|---|---|--|--|------------------|
| client-output-<br>buffer-limit-<br>slave-soft-<br>seconds | When the<br>client-<br>output-<br>buffer-slave-<br>soft-limit<br>parameter is<br>exceeded for<br>more than<br>the value of<br>this<br>parameter,<br>the server<br>drops the<br>connection.<br>The smaller<br>the value, the<br>easier the<br>disconnection. | Single-node<br>instances do<br>not have this<br>parameter. | 0–60<br>Unit: second                   | 60               |
| client-output-<br>buffer-slave-<br>hard-limit             | Hard limit on<br>the output<br>buffer of<br>replica clients.<br>Once the<br>output buffer<br>exceeds the<br>hard limit, the<br>client is<br>immediately<br>disconnected.<br>The smaller<br>the value, the<br>easier the<br>disconnection.                   | Single-node<br>instances do<br>not have this<br>parameter. | 0–<br>17,179,869,18<br>4<br>Unit: byte | 1,717,986,918    |
| Parameter                                     | Description  | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|---|--|--|--|------------------|
| client-output-<br>buffer-slave-<br>soft-limit | Soft limit on<br>the output<br>buffer of<br>replica clients.<br>Once the<br>output buffer<br>exceeds the<br>soft limit and<br>continuously<br>remains<br>above the<br>limit for the<br>time specified<br>by the client-<br>output-<br>buffer-limit-<br>slave-soft-<br>seconds<br>parameter,<br>the client is<br>disconnected.<br>The smaller<br>the value, the<br>easier the<br>disconnection. | Single-node<br>instances do<br>not have this<br>parameter. | 0–<br>17,179,869,18<br>4<br>Unit: byte | 1,717,986,918    |

| Parameter   | Description  | Exception<br>Scenario   | Value Range   | Default<br>Value   |
|---|--|---|---|--|
| maxmemory-<br>policy The policy<br>applied when<br>the<br>maxmemory<br>limit is<br>reached. 8<br>values are<br>available. | The policy<br>applied when<br>the<br>maxmemory<br>limit is<br>reached. 8<br>values are<br>available. | - • • • • • • • • • • • • • • • • • • •   | <ul> <li>volatile-<br/>lru: Evict<br/>keys by<br/>trying to<br/>remove the<br/>less<br/>recently<br/>used (LRU)<br/>keys first,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>lru: Evict<br/>keys by<br/>trying to<br/>remove the<br/>LRU keys<br/>first.</li> </ul> | volatile-lru<br>NOTE<br>If the DCS<br>Redis instance<br>is created<br>before July<br>2020 and this<br>parameter has<br>not been<br>modified, the<br>default value<br>is <b>noeviction</b> .<br>If the instance<br>is created<br>after July<br>2020, the<br>default value<br>is <b>volatile-lru</b> . |
|   |  | <ul> <li>volatile-<br/>random:<br/>Evict keys<br/>randomly,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> </ul> |   |  |
|   |  |   | • allkeys-<br>random:<br>Evict keys<br>randomly.  |  |
|   |  |   | • volatile-<br>ttl: Evict<br>keys with<br>an expire<br>set, and try<br>to evict<br>keys with a<br>shorter<br>time to live<br>(TTL) first.   |  |
|   |  |   | <ul> <li>noeviction         <ul> <li>Do not</li> <li>delete any</li> <li>keys and</li> </ul> </li> </ul>  |  |

| Parameter            | Description   | Exception<br>Scenario  | Value Range   | Default<br>Value |
|----------------------|---|--|---|------------------|
|                      |   |  | only return<br>errors<br>when the<br>memory<br>limit was<br>reached.  |                  |
|                      |   |  | <ul> <li>volatile-<br/>lfu: Evict<br/>keys by<br/>trying to<br/>remove the<br/>less<br/>frequently<br/>used (LFU)<br/>keys first,<br/>but only<br/>among<br/>keys that<br/>have an<br/>expire set.</li> <li>allkeys-<br/>lfu: Evict<br/>keys by<br/>trying to<br/>remove the<br/>LFU keys<br/>first.</li> <li>For details<br/>about eviction<br/>policies, see<br/>the Redis<br/>official<br/>website.</li> </ul> |                  |
| lua-time-limit       | Maximum<br>time allowed<br>for executing<br>a Lua script.                         | -  | 100–5,000<br>Unit:<br>millisecond   | 5,000            |
| master-read-<br>only | Sets the<br>instance to be<br>read-only. All<br>write<br>operations<br>will fail. | Proxy Cluster<br>instances do<br>not have this<br>parameter. | <ul><li>yes</li><li>no</li></ul>  | no               |

| Parameter  | Description  | Exception<br>Scenario  | Value Range | Default<br>Value |
|------------|--|--|-------------|------------------|
| maxclients | The maximum number of clients allowed to be concurrently connected to a DCS instance. The larger the value, the more costly the connection to the server, which affects the server performance and increases the command latency. An excessively small value may constrain the server performance. This parameter specifies the maximum number of connections on a single node (single shard). | Read/Write<br>splitting<br>instances do<br>not support<br>this<br>parameter. | 1000-50,000 | 10,000           |

| Parameter              | Description  | Exception<br>Scenario | Value Range                             | Default<br>Value |
|------------------------|--|-----------------------|---|------------------|
|                        | master/<br>standby:<br>Maximum<br>connection<br>s on a<br>single<br>node =<br>Maximum<br>connection<br>s of the<br>instance  |                       |   |                  |
| proto-max-<br>bulk-len | Maximum<br>size of a<br>single<br>element<br>request. Set<br>this<br>parameter to<br>be greater<br>than the<br>customer<br>request<br>length.<br>Otherwise,<br>the request<br>cannot be<br>executed. | -                     | 1,048,576–<br>536,870,912<br>Unit: byte | 536,870,912      |

| Parameter             | Description   | Exception<br>Scenario                                      | Value Range                            | Default<br>Value |
|-----------------------|---|--|--|------------------|
| repl-backlog-<br>size | The<br>replication<br>backlog size.<br>The backlog is<br>a buffer that<br>accumulates<br>replica data<br>when replicas<br>are<br>disconnected<br>from the<br>master. When<br>a replica<br>reconnects, a<br>partial<br>synchronizatio<br>n is<br>performed to<br>synchronize<br>the data that<br>was missed<br>while replicas<br>were<br>disconnected. | -  | 16,384–<br>1,073,741,824<br>Unit: byte | 1,048,576        |
| repl-backlog-<br>ttl  | The amount<br>of time, in<br>seconds,<br>before the<br>backlog<br>buffer is<br>released,<br>starting from<br>the last a<br>replica was<br>disconnected.<br>The value <b>0</b><br>indicates that<br>the backlog is<br>never<br>released.   | -  | 0–604,800<br>Unit: second              | 3,600            |
| repl-timeout          | Replication<br>timeout.   | Single-node<br>instances do<br>not have this<br>parameter. | 30–3,600<br>Unit: second               | 60               |

| Parameter                    | Description  | Exception<br>Scenario | Value Range | Default<br>Value |
|------------------------------|--|-----------------------|-------------|------------------|
| hash-max-<br>ziplist-entries | The<br>maximum<br>number of<br>hashes that<br>can be<br>encoded<br>using ziplist, a<br>data structure<br>optimized to<br>reduce<br>memory use.   | -                     | 1–10,000    | 512              |
| hash-max-<br>ziplist-value   | The largest<br>value allowed<br>for a hash<br>encoded<br>using ziplist, a<br>special data<br>structure<br>optimized for<br>memory use.   | -                     | 1–10,000    | 64               |
| set-max-<br>intset-entries   | When a set is<br>composed<br>entirely of<br>strings and<br>number of<br>integer<br>elements is<br>less than this<br>parameter<br>value, the set<br>is encoded<br>using intset, a<br>data structure<br>optimized for<br>memory use. | -                     | 1–10,000    | 512              |
| zset-max-<br>ziplist-entries | The<br>maximum<br>number of<br>sorted sets<br>that can be<br>encoded<br>using ziplist, a<br>data structure<br>optimized to<br>reduce<br>memory use.  | -                     | 1–10,000    | 128              |

| Parameter                  | Description  | Exception<br>Scenario | Value Range | Default<br>Value |
|----------------------------|--|-----------------------|-------------|------------------|
| zset-max-<br>ziplist-value | The largest<br>value allowed<br>for a sorted<br>set encoded<br>using ziplist, a<br>special data<br>structure<br>optimized for<br>memory use. | -                     | 1–10,000    | 64               |

| Parameter                         | Description  | Exception<br>Scenario  | Value Range                          | Default<br>Value |
|-----------------------------------|--|--|--------------------------------------|------------------|
| latency-<br>monitor-<br>threshold | The minimum<br>amount of<br>latency that<br>will be logged<br>as latency<br>spikes<br>If this<br>parameter is<br>set to <b>0</b> ,<br>latency<br>monitoring is<br>disabled. If<br>this<br>parameter is<br>set to a value<br>greater than<br>0, all events<br>blocking the<br>server for a<br>time greater<br>than the<br>configured<br>value will be<br>logged.<br>To obtain<br>statistics data,<br>and configure<br>and enable<br>latency<br>monitoring,<br>run the<br><b>LATENCY</b><br>command. | Proxy Cluster<br>instances do<br>not have this<br>parameter. | 0-86,400,000<br>Unit:<br>millisecond | 0                |

| Parameter | Description   | Exception<br>Scenario | Value Range | Default<br>Value |
|-----------|---|-----------------------|-------------|------------------|
|           | NOTE<br>The latency-<br>monitor-<br>threshold<br>parameter is<br>usually used<br>for fault<br>location. After<br>locating faults<br>based on the<br>latency<br>information<br>collected,<br>change the<br>value of<br>latency-<br>monitor-<br>threshold to<br>0 to avoid<br>unnecessary<br>latency. |                       |             |                  |

| Parameter                      | Description  | Exception<br>Scenario  | Value Range   | Default<br>Value |
|--------------------------------|--|--|---|------------------|
| notify-<br>keyspace-<br>events | Controls<br>which<br>keyspace<br>events<br>notifications<br>are enabled<br>for. If this<br>parameter is<br>configured,<br>the Redis<br>Pub/Sub<br>feature will<br>allow clients<br>to receive an<br>event<br>notification<br>when a Redis<br>data set is<br>modified. This<br>parameter is<br>disabled only<br>when it is left<br>blank. | Proxy Cluster<br>instances do<br>not have this<br>parameter. | A<br>combination<br>of different<br>values can be<br>used to<br>enable<br>notifications<br>for multiple<br>event types.<br>Possible<br>values<br>include:<br>K: Keyspace<br>events,<br>published<br>with the<br>keyspace@*<br>prefix<br>E: Keyevent<br>events,<br>published<br>with<br>keyevent@*<br>prefix<br>g: Generic<br>commands<br>(non-type<br>specific) such<br>as DEL,<br>EXPIRE, and<br>RENAME<br>\$: String<br>commands<br>I: List<br>commands<br>S: Set<br>commands<br>s: Set<br>commands<br>x: Expired<br>events<br>(events<br>generated<br>every time a<br>key expires) | Ex               |

| Parameter | Description | Exception<br>Scenario | Value Range   | Default<br>Value |
|-----------|-------------|-----------------------|---|------------------|
|           |             |                       | e: Evicted<br>events<br>(events<br>generated<br>when a key is<br>evicted from<br>maxmemory)   |                  |
|           |             |                       | A: an alias for<br>"g\$lshzxe"  |                  |
|           |             |                       | The<br>parameter<br>value must<br>contain either<br>K or E. A<br>cannot be<br>used together<br>with any of<br>the characters<br>in "g\$lshzxe".<br>For example,<br>the value Kl<br>means that<br>Redis will<br>notify<br>Pub/Sub<br>clients about<br>keyspace<br>events and<br>list<br>commands.<br>The value<br>AKE means<br>Redis will<br>notify<br>Pub/Sub<br>clients about<br>all events. |                  |

| Parameter                             | Description  | Exception<br>Scenario  | Value Range                                  | Default<br>Value |
|---------------------------------------|--|--|--|------------------|
| slowlog-log-<br>slower-than           | Slow queries<br>cover<br>scheduled<br>commands<br>whose<br>execution is<br>delayed.<br><b>slowlog-log-</b><br><b>slower-than</b><br>is the<br>maximum<br>time allowed<br>for command<br>execution. If<br>this threshold<br>is exceeded,<br>Redis will<br>record the<br>query.                                    | -  | 0–1,000,000<br>Unit:<br>microsecond          | 10,000           |
| proxy-<br>slowlog-log-<br>slower-than | Slow queries<br>of a proxy<br>cover<br>scheduled<br>commands<br>whose<br>execution is<br>delayed.<br><b>proxy-</b><br><b>slowlog-log-</b><br><b>slower-than</b><br>is the<br>maximum<br>time allowed<br>for command<br>execution. If<br>this threshold<br>is exceeded,<br>the proxy will<br>record the<br>query. | Currently,<br>only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances in<br>the CN East-<br>Shanghai2<br>and CN<br>South-<br>Guangzhou<br>regions. | 30,000–<br>2,000,000<br>Unit:<br>microsecond | 256,000          |

| Parameter                     | Description  | Exception<br>Scenario  | Value Range | Default<br>Value |
|-------------------------------|--|--|-------------|------------------|
| slowlog-max-<br>len           | The<br>maximum<br>allowed<br>number of<br>slow queries<br>that can be<br>logged. Slow<br>query log<br>consumes<br>memory, but<br>you can<br>reclaim this<br>memory by<br>running the<br><b>SLOWLOG</b><br><b>RESET</b><br>command.               | -  | 0-1000      | 128              |
| proxy-<br>slowlog-max-<br>len | The<br>maximum<br>allowed<br>number of<br>slow queries<br>of a proxy<br>that can be<br>logged. Slow<br>query log<br>consumes<br>memory, but<br>you can<br>reclaim this<br>memory by<br>running the<br><b>SLOWLOG</b><br><b>RESET</b><br>command. | Currently,<br>only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances in<br>the CN East-<br>Shanghai2<br>and CN<br>South-<br>Guangzhou<br>regions. | 0-1000      | 128              |

| Parameter                             | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---------------------------------------|---|--|--|------------------|
| auto-kill-<br>timeout-lua-<br>process | When this<br>parameter is<br>enabled,<br>processes<br>running the<br>lua script are<br>killed when<br>their<br>execution<br>times out.<br>However,<br>scripts with<br>write<br>operations are<br>not killed, but<br>their nodes<br>automatically<br>restart (if<br>persistence<br>has been<br>enabled for<br>the instance)<br>without<br>saving the<br>write<br>operations. | Single-node<br>instances and<br>DCS Redis 3.0<br>instances do<br>not have this<br>parameter. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul> | no               |

| Parameter                               | Description   | Exception<br>Scenario   | Value Range  | Default<br>Value |
|---|---|---|--|------------------|
| audit-log-<br>customer-<br>command-list | Commands to<br>record in<br>audit logs<br>(only write<br>commands<br>are recorded<br>by default.)<br>This<br>parameter is<br>valid only<br>when the<br>audit log<br>function is<br>enabled. | Viewing<br>Audit Logs of<br>a DCS Redis<br>Instance is<br>available only<br>in certain<br>regions. This<br>parameter is<br>displayed only<br>for Proxy<br>Cluster<br>instances<br>when the<br>audit log<br>feature is<br>supported. | A maximum<br>of 10<br>commands<br>are allowed.<br>For each<br>command,<br>use up to 10<br>characters<br>including<br>letters,<br>periods (.),<br>hyphens (-),<br>and<br>underscores<br>(_), and start<br>and end with<br>a letter.<br>Separate<br>multiple<br>commands<br>with spaces,<br>and end your<br>input with a<br>space. |                  |

| Parameter               | Description   | Exception<br>Scenario   | Value Range  | Default<br>Value |
|-------------------------|---|---|--|------------------|
| backend-<br>master-only | Read/Write<br>splitting is<br>disabled by<br>default for<br>Proxy Cluster<br>instances. In<br>this case, read<br>and write<br>requests are<br>allocated to<br>the master<br>node of a<br>Proxy Cluster<br>instance.<br>For Proxy<br>Cluster<br>instances with<br>read/write<br>splitting<br>enabled, read<br>requests are<br>allocated<br>evenly to<br>each master<br>or standby<br>node of a<br>Proxy Cluster<br>instance while<br>write requests<br>are allocated<br>to the master<br>node. | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter. | <ul> <li>yes:<br/>disables<br/>read/write<br/>splitting.</li> <li>no:<br/>enables<br/>read/write<br/>splitting.</li> </ul> | yes              |

| Parameter                             | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---------------------------------------|---|--|--|------------------|
| read-only-<br>slave-when-<br>wr-split | By default,<br>Proxy Cluster<br>instances<br>perform read<br>operations<br>only on<br>standby<br>nodes. To<br>perform read<br>operations on<br>both master<br>and standby<br>nodes,<br>configure this<br>parameter. | Only Proxy<br>Cluster and<br>read/write<br>splitting<br>instances<br>have this<br>parameter<br>and it is only<br>applicable<br>when read/<br>write splitting<br>is enabled.<br>Read/Write<br>splitting is<br>enabled by<br>default for<br>read/write<br>splitting<br>instances. For<br>Proxy Cluster<br>instances, it<br>can be<br>enabled by<br>setting<br>parameter<br><b>backend-</b><br><b>master-only</b><br>to <b>no</b> . | <b>yes</b> : Read<br>only on<br>standby<br>nodes.<br><b>no</b> : Read on<br>both master<br>and standby<br>nodes. | yes              |

| Parameter                                   | Description   | Exception<br>Scenario  | Value Range  | Default<br>Value |
|---|---|--|--|------------------|
| dispatch-<br>pubsub-to-<br>fixed-shard      | This<br>parameter<br>specifies<br>whether<br>pub/sub<br>channels are<br>on the shard<br>of slot 0.<br>When this<br>parameter is<br>enabled, the<br>pub/sub<br>processing<br>logic is<br>consistent<br>with that of<br>single-node<br>instances. You<br>are advised to<br>enable this<br>parameter if<br>you do not<br>depend<br>heavily on<br>pub/sub. If<br>you depend<br>heavily on<br>pub/sub, use<br>the default<br>configuration<br>to allocate<br>subscriptions<br>to all shards. | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter.            | <ul> <li>yes: Enable this parameter to allocate subscriptio n channels to the shard of slot 0.</li> <li>no: Disable this parameter to allocate channels to the shard of each channel-hashed slot.</li> </ul> | no               |
| readonly-lua-<br>route-to-<br>slave-enabled | If enabled,<br>read-only Lua<br>scripts of<br>read-only<br>users are<br>executed and<br>routed to the<br>standby node.  | Only read/<br>write splitting<br>instances<br>support this<br>parameter. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul>   | no               |
| cluster-<br>sentinel-<br>enabled            | ster-<br>ntinel-<br>abled the instance.   |  | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul>   | no               |

| Parameter                 | Description   | Exception<br>Scenario   | Value Range  | Default<br>Value |
|---------------------------|---|---|--|------------------|
| scan-support-<br>wr-split | The <b>SCAN</b><br>command is<br>executed on<br>the master<br>node when<br>this<br>parameter is<br>disabled, or is<br>executed on<br>the standby<br>node<br>otherwise.<br>Enabling this<br>parameter<br>relieves SCAN<br>commands on<br>the master<br>node. But<br>newly written<br>data in the<br>master node<br>may not be<br>synchronized<br>to replicas in<br>time. | Only Proxy<br>Cluster<br>instances<br>have this<br>parameter.<br>Proxy Cluster<br>instances<br>created<br>earlier may<br>not support<br>this<br>parameter. In<br>this case,<br>contact<br>customer<br>service to<br>upgrade<br>instances. | <ul> <li>yes:<br/>enabled</li> <li>no:<br/>disabled</li> </ul> | no               |

#### D NOTE

- The default values and value ranges of the maxclients, reserved-memory-percent, client-output-buffer-slave-soft-limit, and client-output-buffer-slave-hard-limit parameters are related to the instance specifications. Therefore, these parameters are not displayed in the parameter template.
- 2. For more information about the parameters described in Table 6-5, visit https:// redis.io/topics/memory-optimization.

#### Step 8 Click OK.

----End

#### Modifying or Deleting Custom Templates

- **Step 1** Log in to the DCS console.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- Step 3 In the navigation pane, choose Parameter Templates.
- Step 4 Choose the Custom Templates tab.

- **Step 5** To edit a custom parameter template, use either of the following ways:
  - Locate the row containing the desired template and click **Edit** in the **Operation** column.
    - a. Change the name or modify the description.
    - b. In the **Parameters** area, select **Modifiable parameters**. In the row containing the desired parameter, enter a value in the **Assigned Value** column. **Table 6-6** describes the parameters. In most cases, retain the default values.
    - c. Click **OK**.
  - Click the name of a custom template.
    - a. Select **Modifiable parameters**. Enter a keyword in the search box to search for a parameter by its name.
    - b. Click Modify.
    - c. In the row containing the desired parameter, enter a value in the Assigned Value column. Table 6-6 describes the parameters. In most cases, retain the default values.
    - d. Click Save.
- **Step 6** To delete custom templates, click **Delete** in the **Operation** column on the right of the templates to be deleted.

Click Yes.

----End

## 6.6 Configuring DCS Instance Tags

Tags facilitate DCS instance identification and management. Tags can be added in instance creation, or on the instance details page. A maximum of 20 tags are allowed for a DCS instance.

#### **NOTE**

If your organization has configured tag policies for DCS, add tags to DCS instances based on the tag policies. If a tag does not comply with the policies, tag addition may fail. Contact your organization administrator to learn more about tag policies.

A tag consists of a tag key and tag value. **Table 6-7** describes the naming rules for them.

| Parameter | Requirements  |
|-----------|---|
| Tag key   | Cannot be left blank.   |
|           | <ul> <li>Must be unique for the same instance.</li> </ul>   |
|           | <ul> <li>Consists of a maximum of 128 characters.</li> </ul>  |
|           | <ul> <li>Can contain letters of any<br/>language, digits, spaces, and special<br/>characters : = + - @</li> </ul>   |
|           | • Cannot start or end with a space.   |
|           | • Cannot start with _ <b>sys</b>  |
| Tag value | • Consists of a maximum of 255 characters.  |
|           | <ul> <li>Can contain letters of any<br/>language, digits, spaces, and special<br/>characters : / = + - @</li> </ul> |
|           | • Cannot start or end with a space.   |

 Table 6-7 Tag key and value requirements

#### **Configuring Instance Tags**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of the desired DCS instance to go to the details page.
- **Step 5** Choose **Instance Configuration** > **Tags**.
- **Step 6** Perform the following operations as required:
  - Add a tag
    - a. Click Add/Edit Tag.

If you have created predefined tags, select a predefined pair of tag key and value. To view or create predefined tags, click **View predefined tags**. Then you will be directed to the TMS console.

You can also create new tags by specifying **Tag key** and **Tag value**.

- b. Click OK.
- Modify a tag

Click **Add/Edit Tag**. In the displayed **Add/Edit Tag** dialog box, delete the desired key, add the key again, enter a new tag value, and click **Add**.

• Delete a tag

In the row that contains the desired tag, click **Delete**. In the displayed dialog box, click **Yes**.

----End

## 6.7 Renaming Critical Commands for DCS Instances

Certain high-risk commands can be modified for DCS Redis instances. Once a command is modified, it is only known to the modifier. Running the original command by other users is blocked.

Currently, you can only rename the **COMMAND**, **KEYS**, **FLUSHDB**, **FLUSHALL**, **HGETALL**, **SCAN**, **HSCAN**, **SSCAN**, and **ZSCAN** commands. For Proxy Cluster instances, you can also rename the **DBSIZE** and **DBSTATS** commands.

#### Notes and Constraints

- Only DCS Redis 4.0 and later instances support command renaming.
- An instance will restart when its commands are renamed. Exercise caution.
- Renaming takes effect immediately once it is complete. Renamed commands will not be displayed on the console for security purposes.
- A command can be renamed multiple times. Each new name overwrites the previous name. To restore a high-risk command, or if a renamed command is forgotten, rename the command.
- A command cannot be renamed to other original commands. For example, **KEYS** can be renamed to **KEYS** or **ABC123**, but cannot be renamed to **SCAN**.
- Renaming a command starts only with a letter and contains 4–64 characters of letters, digits, hyphens (-), and underscores (\_).

#### Procedure

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** In the **Operation** column of an instance, choose **More** > **Command Renaming**.
- **Step 5** Select a command, enter a new name, and click **OK**.

In the **Command Renaming** dialog box, click **Add Command** to rename multiple commands at the same time. After renaming commands, you can view the renaming operation record on the **Background Tasks** page.

#### Figure 6-3 Renaming commands

| No. \ominus 🛛 Task Name 🕀 | ID 🕀 | Username 🕀 | Status 🕀   | Start Time \ominus              | End Time \ominus                | Detailed Information $ \Leftrightarrow $ | Operation |
|---------------------------|------|------------|------------|---------------------------------|---------------------------------|--|-----------|
| 1 Command Renaming R      |      |            | Successful | Dec 09, 2024 15:16:24 GMT+08:00 | Dec 09, 2024 15:16:38 GMT+08:00 | Before: command                          | Delete    |

----End

# 6.8 Returning the Real IP Addresses of a Client to DCS (IP Pass-through)

When a client of a DCS Redis 4.0 or later instance connects to the server through a VPC endpoint, the source IP address displayed on the server belongs to the VPC endpoint (starting with 198.19), and it is not the real client IP address.

After the **client IP pass-through** function is enabled, the real IP address and port of a client can be returned when O&M personnel **manage sessions** or run commands such as **Client List**, **Monitor**, and **Slowlog Get**.

For Redis 3.0, run **CLIENT LIST** to view the real client IP address.

#### **Notes and Constraints**

- Currently, the client IP pass-through function can be enabled on the console in the CN North-Beijing1, CN East-Shanghai2, and CN South-Guangzhou regions. To enable client IP pass-through in other regions, contact customer service.
- **SSL encryption** and client IP pass-through cannot be enabled at the same time for instances.
- Client IP pass-through works only on new connections upon sending a Redis command. IP addresses of existing connections still start with 198.19.

#### **Configuring Client IP Pass-through**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** On the **Cache Manager** page, click the name of the DCS instance for which you want to enable client IP pass-through.
- **Step 5** In the **Connection** area, click *level* next to **Client IP Pass-through**.

#### Figure 6-4 Configuring client IP pass-through

| Connection ③           |        |
|------------------------|--------|
| Password Protected     | No     |
| Connection Address     |        |
| Read-only Address ⑦    |        |
| IP Address             |        |
| Public Access          |        |
| Client IP Pass-through | On 🖉 📀 |
| •                      |        |

**Step 6** View the client IP address. (The following uses the **Client List** command as an example.)

In the record that contains "network=vpc", the value of **addr** is the client IP address.

Figure 6-5 Before enabling client IP pass-through



#### **NOTE**

Client IP pass-through works only on new connections upon sending a Redis command. IP addresses of existing connections still start with 198.19.

----End

# 6.9 Exporting a DCS Instance List

A list of instance information can be exported and downloaded in Excel on the DCS console. You can view or compare DCS instance information.

#### Procedure

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

- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- Step 3 In the navigation pane, choose Cache Manager.
- **Step 4** Click **Export** to export all instances by default. To export some instances, select them and click **Export**.
- **Step 5** The **Tasks** page is displayed. When the **Export DCS instance list** task is in the **Successful** state, click **Download** on the right to download the list.

#### Figure 6-7 Exported DCS instance list

| Name     | ID         | Status  | AZ  | Cache | EngInstance   | SpecificsUsed/AvaiConnecticCreated | Billing NVPC  | VPC ID      | Enterpris | e Project |
|----------|------------|---------|-----|-------|---------------|------------------------------------|---------------|-------------|-----------|-----------|
| dcs-trpt | 5e4f4c58-  | Running | AZ1 | Redis | 5. (Single-no | 0.125 0/128 (0%198.19.32May 24,    | 2Free nul     | l null      | default   |           |
| dcs-APIT | e693491b0- | Running |     | Redis | 3.(Master/St  | 2 2/1,536 (172.16.14May 06,        | 2Yearly/Monul | .1 52267da0 | -default  |           |

----End

# 6.10 Performing a Master/Standby Switchover for a DCS Instance

On the DCS console, you can manually switch the master and standby nodes of a master/standby or read/write splitting DCS instance. This operation is used for

special purposes, for example, releasing all service connections or terminating ongoing service operations.

This operation is not available for cluster instances. To perform a manual switchover for a Proxy Cluster or Redis Cluster instance, go to the **Node Management** or **Shards and Replicas** page of the instance. For details, see **Managing DCS Instance Shards and Replicas**.

The instance IP address does not change after a master/standby switchover, so the client does not need to change the connection address.

#### Notes and Constraints

- The instance must be in the **Running** state.
- During a master/standby switchover, services will be interrupted for up to 10 seconds. Before performing this operation, ensure that your application supports connection re-establishment in case of disconnection.
- During a master/standby node switchover, a large amount of resources will be consumed for data synchronization between the master and standby nodes. You are advised to perform this operation during off-peak hours.
- Data of the maser and standby nodes is synchronized asynchronously. Therefore, a small amount of data that is being operated on during the switchover may be lost.

#### Performing a Master/Standby Switchover for a DCS Instance

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click <sup>(Q)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- Step 4 In the Operation column of the instance, choose More > Master/Standby Switchover.
- **Step 5** When the master/standby switchover task is in the **Successful** state, the master/ standby is switched over.

----End

# 6.11 Managing DCS Instance Shards and Replicas

This section describes how to query the shards and replicas of a master/standby, cluster, or read/write splitting DCS Redis instance, and how to manually promote a replica to master.

- By default, a master/standby or read/write splitting instance has only one shard with one master and one replica. You can view the sharding information on the Node Management page. To manually switch the master and replica roles, see Performing a Master/Standby Switchover for a DCS Instance.
- On the **Node Management** page: The failover priority can be edited for master/standby instances with multiple replicas. The IP address of a replica can be removed (only when multiple replicas exist). The information returned

when an instance is accessed at read-only domain names excludes the removed IP addresses.

- A Proxy Cluster or Redis Cluster instance has multiple shards. Each shard has one master and one replica. On the **Node Management** page, you can view the sharding information and manually switch the master and replica roles.
- For details about the number of shards for different instance specifications, see **Redis Cluster** and **Proxy Cluster Redis**.
- You can adjust shards of a cluster instance by referring to Modifying DCS Instance Specifications.

#### **Notes and Constraints**

- This feature is called "Node Management" in some regions and "Shard and Replica" in the other regions. Refer to the console for the actual name.
- This feature is supported by DCS Redis 4.0 instances and later.
- For single-node DCS instances, this feature is supported only in regions where **Node Management** is used.

#### **Managing DCS Instance Shards and Replicas**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click an instance.
- Step 5 Click Node Management or Shards and Replicas.

All shards of the instance are displayed by **Shard Name**, **Shard ID**, and **Replicas** of a shard.

Step 6 Click  $\checkmark$  to display all replicas of a shard by Replica IP Address, Node ID, Replica ID, Status, Role, and AZ.

#### Figure 6-8 Node management (cluster instance)

| Proxies    | Redis Servers |                                      |              |                                      |           |  |          |     |                   |   |
|------------|---------------|--------------------------------------|--------------|--------------------------------------|-----------|--|----------|-----|-------------------|---|
|            |               |                                      |              |                                      |           |  |          |     |                   | С |
| Shard      | d Name        |                                      |              | Shard ID                             |           |  | Replicas |     |                   |   |
| ∧ group-   | ∧ group-0     |                                      |              | 9e28044c-099d-4250-a139-16a1bb00b4de |           |  | 2        |     |                   |   |
| Deplice II | Dåddross      | Node ID                              | Daplica ID   |                                      | Status    |  | Bala     | 47  | Oneration         |   |
| replica in | r Auuress     | NOUGHD                               | Replica in   |                                      | status (j |  | NOTE     | AL. | Operation         |   |
| 192.168.2  | 28.12         | ff8080828bbe9180018bc88d85c84075     | 86abc931-10  | ec-4cdc-b0e5-2ffa62a6bafe            | Running   |  | Master   | AZ1 |                   |   |
| 192.168.2  | 27.236        | ff5080828bbe9180018bc88d85d44076     | db4f6f5a-258 | ic-4abe-90d4-115522ee576b            | Running   |  | Replica  | AZ2 | Promote to Master |   |
|            |               |                                      |              |                                      |           |  |          |     |                   |   |
| ✓ group-   | a-1           | c3d82976-7c0c-4b56-be3/-28e17%cf81fd |              |                                      |           |  | 2        |     |                   |   |
| ✓ group-   | a-2           | 9ea7adec-7700-4212-91e9-1577d9666ca3 |              |                                      |           |  | 2        |     |                   |   |

#### Figure 6-9 Node management (master/standby instance)

| Re | dis Servers        |                                  |                           |                 |                       |         |       |                       |                   | С |
|----|--------------------|----------------------------------|---------------------------|-----------------|-----------------------|---------|-------|-----------------------|-------------------|---|
|    | Shard Name         |                                  |                           | Shard ID        |                       |         | Repli | cas                   |                   |   |
|    | ∧ group-0          |                                  |                           | f5185f84-62ba-4 | ca1-aa99-0a97468793c8 |         | 2     |                       |                   |   |
|    | Replica IP Address | Node ID                          | Replica ID                |                 | Status 💮              | Role    | AZ    | Failover Priority (?) | Operation         |   |
|    | 10.0.0.179         | ff8080828bbe9180018bc882a86c400a | 6cecedf0-9654-48be-9ca8-  | 878fd01afefc    | Sunning               | Master  | AZ1   |                       |                   |   |
|    | 10.0.0.186         | ff8080828bbe9180018bc882a878400b | cfa2af74-3d3c-4d9c-8c29-d | 6ed1d9fd7ae     | Running               | Replica | AZ2   | 100 🖉                 | Remove IP Address |   |
|    |                    |                                  |                           |                 |                       |         |       |                       |                   |   |

Figure 6-10 Node management (single-node instance)

| Redis Servers      |                                  |                                      |            |          |        | С   |
|--------------------|----------------------------------|--------------------------------------|------------|----------|--------|-----|
| Shard Name         |                                  | Shard ID                             |            | Replicas |        |     |
| ∧ group-0          |                                  | 7bfa0e38-0b60-45cf-ae21-d63e1840ffd4 |            | 1        |        |     |
| Replica IP Address | Node ID                          | Replica ID                           | Status (?) |          | Role   | AZ  |
| 10.0.36            | ff8080628b610079018b8fbd694b4e2d |                                      | 8 Running  |          | Master | AZ1 |
|                    |                                  |                                      |            |          |        |     |

You can also perform the following operations on replicas:

Cluster

To promote a replica to the master role, expand a shard and click **Promote to Master** in the row that contains a node whose **Role** is **Replica**.

D NOTE

For **Proxy Cluster** instance, the proxy information (IP address, node ID, and name) can be viewed on the **Node Management** > **Proxies** page. Other types of instances do not have the **Proxies** tab page.

- Master/Standby or read/write splitting
  - a. If a master/standby instance has multiple replicas, click **Remove IP** Address in the row containing a read-only replica. After a replica IP address is removed, the read-only domain name will no longer be resolved to the replica IP address.

If a master/standby instance has only one replica, its IP address cannot be removed.

b. If a master/standby or read/write splitting instance has multiple replicas,

click solution in the **Failover Priority** column to change the priority of the replica to be promoted to master.

If the master fails, the replica with the smallest priority number is automatically promoted to master. For multiple replicas that have the same priority, a selection process will be performed. **0** indicates that the replica will never be automatically promoted, **1** indicates the highest priority, and **100** indicates the lowest priority.

• Single-node

A single-node instance has only one replica. You can view its node information on the **Node Management** page.

----End

# **7** Backing Up or Restoring Instance Data

# 7.1 DCS Backup and Restoration Overview

On the DCS console, you can back up and restore DCS instances.

#### **Importance of DCS Instance Backup**

There is a small chance that dirty data could exist in a DCS instance owing to service system exceptions or problems in loading data from persistence files. In addition, some systems demand not only high reliability but also data security, data restoration, and even permanent data storage.

Currently, data in DCS instances can be backed up to OBS. If a DCS instance becomes faulty, data in the instance can be restored from backup so that service continuity is not affected.

#### **Backup Modes**

DCS instances support the following backup modes:

Automated backup

You can create a scheduled backup policy on the DCS console. Then, data in the chosen DCS instances will be automatically backed up at the scheduled time.

You can choose the days of the week on which automated backup will run. Backup data will be retained for a maximum of seven days. Backup data older than seven days will be automatically deleted.

The primary purpose of automated backups is to create complete data replicas of DCS instances so that the instance can be quickly restored if necessary.

Manual backup

Backup requests can be issued manually. Data in the chosen DCS instances will be backed up to OBS.

Before performing high-risk operations, such as system maintenance or upgrade, back up DCS instance data.

#### Impact on DCS Instances During Backup

# Backup tasks are run on standby cache nodes, without incurring any downtime.

In the event of full synchronization of master and standby nodes or heavy instance load, it takes a few minutes to complete data synchronization. If instance backup starts before data synchronization is complete, the backup data will be slightly behind the data in the master cache node.

New data changes on the master node during an ongoing backup are not included in the backup.

#### **Additional Information About Data Backup**

- Instance type
  - Redis: Only master/standby, read/write splitting, Proxy Cluster, and Redis Cluster instances can be backed up and restored, while singlenode instances cannot. You can export data of a single-node instance to an RDB file using redis-cli. For details, see How Do I Export DCS Redis Instance Data?
  - Memcached: Only master/standby instances can be backed up and restored, while single-node instances cannot.
- Backup mechanisms

DCS for Redis 3.0 persists data to AOF files. Basic edition DCS for Redis 4.0 and later persist data to RDB or AOF files in manual backup mode, and to RDB files in automatic backup mode.

To export RDB backup files of DCS Redis 3.0 instances, run the **redis-cli -h** *{redis\_address}* -**p 6379 [-a** *{password}*] --**rdb** *{output.rdb}* command in redis-cli.

**NOTE** 

- DCS for Redis 3.0 is no longer provided. You can use DCS for Redis 4.0 or later instead.
- For a single-node DCS Redis 3.0 instance on which the **SYNC** command can be run, you can run this command to export the RDB file. For a Proxy Cluster DCS Redis 3.0 instance, the **SYNC** command cannot be run due to the architecture. Therefore, the RDB file cannot be exported.

Backup tasks are run on standby cache nodes. DCS instance data is backed up by compressing and storing the data persistence files from the standby cache node to OBS.

DCS checks instance backup policies once an hour. If a backup policy is matched, DCS runs a backup task for the corresponding DCS instance.

• Backup time

Back up instance data during off-peak periods.

• Storage of backup files

Backup files are stored to OBS.

Handling exceptions in automated backup

If an automated backup task is triggered while the DCS instance is restarting or being scaled up, the backup task will be run in the next cycle. If backing up a DCS instance fails or the backup is postponed because another task is in progress, DCS will try to back up the instance in the next cycle. A maximum of three retries are allowed within a single day.

• Retention period of backup data

Automated backup files are retained for up to seven days. You can configure the retention period. At the end of the retention period, most backup files of the DCS instance will be automatically deleted, but at least the most recent backup record will be retained.

The latest backup files (up to 24) are always stored unless they are manually deleted.

#### **NOTE**

- A total of 24 latest backups (automatic and manual) can be stored. To store the 25<sup>th</sup> backup, the earliest one will be automatically deleted.
- Deleting an instance removes its backups. To restore them, download and save them in advance.
- Exercise caution when deleting all backup files.

#### **Restoring Data**

- Data restoration process
  - a. You can initiate a data restoration request using the DCS console.
  - b. DCS obtains the backup file from OBS.
  - c. Read/write to the DCS instance is suspended.
  - d. The original data persistence file of the master cache node is replaced by the backup file.
  - e. The new data persistence file (that is, the backup file) is reloaded.
  - f. Data is restored, and the DCS instance starts to provide read/write service again.
- Impact on service systems

Restoration tasks are run on master cache nodes. During restoration, data cannot be written into or read from instances.

• Handling data restoration exceptions

If a backup file is corrupted, DCS will try to fix the backup file while restoring instance data. If the backup file is successfully fixed, the restoration proceeds. If the backup file cannot be fixed, the master/standby DCS instance will be changed back to the state in which it was before data restoration.

# 7.2 Backing up DCS Instances Automatically

On the DCS console, you can configure an automatic backup policy. The system then backs up data in your instances according to the backup policy.

By default, automatic backup is disabled. To enable it, perform the operations described in this section. Single-node instances do not support backup and restoration.

If automatic backup is not required, disable the automatic backup function in the backup policy.

#### **Notes and Constraints**

Available for master/standby, cluster, and read/write splitting DCS instances that are in the **Running** state.

#### **Configuring Automated Backup Policies**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.

Filter DCS instances to find the desired DCS instance. Currently, you can search instances by name, specification, ID, IP address, AZ, status, instance type, cache engine, and many other attributes.

- **Step 4** Click the name of the desired DCS instance to go to the instance details page.
- **Step 5** On the instance details page, click **Backups & Restorations**.
- **Step 6** Slide **()** to the right to enable automatic backup. Backup policies will be displayed.

| Parameter                  | Description   |  |  |  |
|----------------------------|---|--|--|--|
| Backup Schedule            | Day of a week on which data in the chosen DCS instance is automatically backed up.  |  |  |  |
|                            | You can select one or multiple days of a week.  |  |  |  |
| Retention Period<br>(days) | The number of days that automatically backed up data is retained.   |  |  |  |
|                            | Backup data will be permanently deleted at the end of retention period and cannot be restored. Value range: 1–7.  |  |  |  |
| Start Time                 | Time at which automatic backup starts. Value: the full hour between 00:00 to 23:00  |  |  |  |
|                            | DCS checks backup policies once every hour. If the backup start time in a backup policy has arrived, data in the corresponding instance is backed up.   |  |  |  |
|                            | <b>NOTE</b><br>Instance backup takes 5 to 30 minutes. The data added or modified<br>during the backup process will not be backed up. To reduce the<br>impact of backup on services, it is recommended that data should<br>be backed up during off-peak periods. |  |  |  |
|                            | Only instances in the <b>Running</b> state can be backed up.  |  |  |  |

Table 7-1 Parameters in a backup policy

#### Step 7 Click OK.

The automated backup can be disabled. The backup policies can be modified.

**Step 8** Automatic backup starts at the scheduled time. You can view backup records on the current page.

After the backup is complete, click **Download**, **Restore**, or **Delete** next to the backup record as required.

----End

# 7.3 Backing up DCS Instances Manually

You can manually back up data in DCS instances in a timely manner. This section describes how to manually back up data in master/standby instances using the DCS console.

#### **Notes and Constraints**

- A total of 24 latest backups (automatic and manual) can be stored. To store the 25<sup>th</sup> backup, the earliest one will be automatically deleted.
- Deleting an instance removes its backups. To restore them, download and save them in advance.
- Exercise caution when deleting all backup files.
- Available for master/standby, cluster, and read/write splitting DCS instances that are in the **Running** state.

#### **Backing up DCS Instances Manually**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click <sup>(2)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.

Filter DCS instances to find the desired DCS instance. Currently, you can search instances by name, specification, ID, IP address, AZ, status, instance type, cache engine, and many other attributes.

- **Step 4** Click the name of the desired DCS instance to go to the details page.
- **Step 5** On the instance details page, click **Backups & Restorations**.

#### Step 6 Click Create Backup.

Step 7 Select RDB or AOF for the backup file format.

#### **NOTE**

Basic and professional (performance) edition DCS Redis 4.0 and later instances support backup files in RDB or AOF formats. Professional (storage) edition instances support the RDB format. DCS Redis 3.0 instances support only the RDB format.

If you select AOF, data will be backed up on the standby node first. The standby node's AOF will be rewritten.

#### Step 8 In the Create Backup dialog box, click OK.

Information in the **Description** text box cannot exceed 128 bytes.

#### D NOTE

Instance backup takes 10 to 15 minutes. The data added or modified during the backup process will not be backed up.

----End

# 7.4 Restoring DCS Instances

This section describes how to restore instances on the DCS console. This function helps restore instances deleted by mistake.

To migrate backup data to other DCS instances, see **Backup Import Between DCS Redis Instances**.

#### Notes and Constraints

- You can **enable or disable multi-DB** for a Proxy Cluster instance. Data backed up when multi-DB is enabled cannot be restored to the instance after multi-DB is disabled.
- This function becomes unavailable after the following changes.
  - Instance scale-in
  - Cluster instance scale-out
  - Instance type changes (from master/standby to read/write splitting not included)

#### Prerequisites

- A master/standby, cluster, or read/write splitting DCS instance is in the **Running** state.
- A backup task has been run to back up data in the instance to be restored and the backup task succeeded.

#### Procedure

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.

Filter DCS instances to find the desired DCS instance. Currently, you can search instances by name, specification, ID, IP address, AZ, status, instance type, cache engine, and many other attributes.

- **Step 4** Click the name of the desired DCS instance to go to the instance details page.
- **Step 5** On the instance details page, click **Backups & Restorations**.

A list of historical backup tasks is then displayed.

**Step 6** Click **Restore** in the row containing the chosen backup task.

Step 7 Click OK to start instance restoration.

Information in the **Description** text box cannot exceed 128 bytes.

You can view the results of all restoration tasks on the **Restoration History** page. The records cannot be deleted.

**NOTE** 

- Instance restoration takes 1 to 30 minutes.
- While being restored, DCS instances do not accept data operation requests from clients because existing data is being overwritten by the backup data.

----End

## 7.5 Downloading DCS Instance Backup Files

Automatically backed up data can be retained for a maximum of 7 days. Manually backed up data is not free of charge and takes space in OBS. Due to these limitations, you are advised to download the RDB and AOF backup files and permanently save them on the local host.

This function is supported only by master/standby, read/write splitting, and cluster instances, and not by single-node instances. To export the data of a single-node instance to an RDB file, you can use redis-cli. For details, see **How Do I Export DCS Redis Instance Data?** 

To export the data of a master/standby, read/write splitting, or cluster instance, do as follows:

 Redis 3.0: Export the instance data to AOF files by using the DCS console, or to RDB files by running the redis-cli -h {redis\_address} -p 6379 -a {password} --rdb {output.rdb} command by using redis-cli.

**NOTE** 

DCS for Redis 3.0 is no longer provided. You can use DCS for Redis 4.0 or later instead.

- Redis 4.0 and later: Export the instance data to AOF or RDB files by using the DCS console.
- For a cluster instance with multiple shards, the downloaded backup contains multiple files for each shard.

#### Prerequisites

The instance has been backed up and the backup is still valid.

#### **Downloading DCS Instance Backup Files**

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

- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- Step 3 In the navigation pane, choose Cache Manager.

Filter DCS instances to find the desired DCS instance.

- **Step 4** Click the name of the DCS instance to display more details about the DCS instance.
- **Step 5** On the instance details page, click **Backups & Restorations**.

A list of historical backup tasks is then displayed.

- **Step 6** Click **Download** in the row containing the chosen backup task.
- **Step 7** Download the backup by URL or OBS.
  - By URL

#### Figure 7-1 By URL

| Download Backup File   |   |            | × |  |  |  |  |  |  |
|--|---|------------|---|--|--|--|--|--|--|
| By URL By OBS  |   |            |   |  |  |  |  |  |  |
| 1. After you set the URL validity pe<br>URLs expire. Properly keep the   | 1. After you set the URL validity period and click Query, the system generates temporary backup file URLs. Download the backup files before the URLs expire. Properly keep the URLs to avoid data breach. |            |   |  |  |  |  |  |  |
| Step 1: Set the URL validity period and click Query.         URL Validity Period       0       hours       30       minutes         The URL validity period must be 5 minutes to 12 hours.       Query         Step 2: Download the backup file in the URL list. Copy all URLs         AUUDI a will write at Sec 24. 2024 15:20-25 CML 09:00. After the entries and entries to activity of the activity of the URL is a second of the se |   |            |   |  |  |  |  |  |  |
| File   | URL   | Operation  |   |  |  |  |  |  |  |
| 20240924145832_redis.rdb   | https://bucket15031547a54245  | a Download |   |  |  |  |  |  |  |
| Close  |   |            |   |  |  |  |  |  |  |

a. Set the URL validity period and click **Query**. A temporary backup file URL will be generated. The URL has a validity period, and needs to be generated again after that.

URL validity period: 5 minutes to 12 hours.

b. In the URL list, click **Download** to download files.

If your account has enabled **critical operation protection**, downloading the backup file requires authentication, as shown in **Figure 7-3**. If not, the backup file will be downloaded to the local.

D NOTE

- Click Copy all URLs or the copy icon after a URL to copy URLs.
- If you choose to copy URLs, use quotation marks to quote the URLs when running the wget command in Linux. For example:

wget 'https://obsEndpoint.com:443/redisdemo.rdb? parm01=value01&parm02=value02'

This is because the URL contains the special character and (&), which will confuse the **wget** command. Quoting the URL facilitates URL identification.

- Keep the backup files and URLs secure to prevent data leakage.
- By OBS
Figure 7-2 By OBS

| By URL By OBS  |                                |  |  |
|--|--------------------------------|--|--|
| -(1)<br>Prepare  | Attach External Bucket         | Download Backup File                                   |  |
| Install OBS Browser+ and log in as<br>user . For details on how<br>to log in to OBS Browser+, see OBS<br>Help. | Bucket:<br>bucket;<br>f3186c 🗇 | Backup file:<br>b0cdd<br>06/20240712093305_redis.rdb 🗇 |  |
|  | Close                          |  |  |

- a. Click By OBS, click OBS Browser+ in the Prepare area.
- b. Install the OBS Browser+ client. Log in to OBS Browser+ using the Huawei Cloud account by referring to Logging In to OBS Browser+.
- c. Add an external bucket to the OBS Browser+ client. For details, see Adding an External Bucket.

Use the bucket name in the Attach External Bucket area.

d. Click the bucket name. Search for backup files in the bucket. For details, see **Searching for a File or Folder**.

The backup file path in the **Download Backup File** area contains the name of the backup file and folder.

e. Click the download icon on the right to download the backup file.

----End

#### (Optional) Enabling Critical Operation Protection

Downloading backup files is a critical operation that can be protected. After operation protection is enabled, a verification code will be needed to download backup files, improving data security.

#### D NOTE

- An administrator can configure critical operation protection, and IAM users can only view the configurations. If an IAM user needs to modify the configurations, the user can request the administrator to perform the modification or grant the required permissions.
- For more information about critical operations, see Critical Operation Protection.
- **Step 1** On the management console, hover over the username in the upper right corner, and choose **Security Settings** from the drop-down list.
- **Step 2** On the **Security Settings** page, click the **Critical Operations** tab.
- Step 3 Click Enable next to Operation Protection to enable operation protection.
- **Step 4** After operation protection is enabled, the **Confirm** dialog box is displayed when you download DCS backup files. You can download backup files only after identity authentication is complete.

#### Figure 7-3 Confirm

| Confirm  |   | × |
|--|---|---|
| You have enabled operation protection. A verification code is required for<br>this operation. To disable operation protection, click Disable next to<br>Operation Protection on the Critical Operations tab page on the Security<br>Settings page. Disable Identity Verification | × |   |
| Verification Method 🔷 SMS 💿 Email 🔷 Virtual MFA device 🕐   |   |   |
| Email s .com C Change  |   |   |
| Verification Code 8-digit code Send Code   |   |   |
| OK Cancel  |   |   |

----End

# 8 Changing an Instance

#### 8.1 Modifying DCS Instance Specifications

On the DCS console, you can change DCS Redis or Memcached instance specifications including the instance type, memory, and replica quantity.

Modifying instance specifications does not affect the connection address, password, security group, and whitelist configurations of the instance. You do not need to restart the instance.

Changing specifications only affects single-node instance data.

#### **Notes and Constraints**

- **Modify instance specifications during off-peak hours.** If the modification failed in peak hours (for example, when memory or CPU usage is over 90% or write traffic surges), try again during off-peak hours.
- If your DCS instances are too old to support specification modification, contact customer service to upgrade the instances.
- Cluster DCS Redis 3.0 instances cannot be vertically scaled.
- If the reserved memory of a DCS Redis 3.0 or Memcached instance is insufficient, the modification may fail when the memory is used up. For details, see **Reserved Memory**.
- Change the replica quantity and capacity separately.
- Only one replica can be deleted per operation.

#### Change of the Instance Type

| Version   | Supported Type<br>Change                    | Precautions  |
|-----------|---|--|
| Redis 3.0 | From single-<br>node to master/<br>standby  | The instance cannot be connected for several seconds and remains read-only for about one minute.   |
|           | From master/<br>standby to Proxy<br>Cluster | <ol> <li>If the data of a master/standby DCS<br/>Redis 3.0 instance is stored in<br/>multiple databases, or in non-DB0<br/>databases, the instance cannot be<br/>changed to the Proxy Cluster type. A<br/>master/standby instance can be<br/>changed to the Proxy Cluster type<br/>only if its data is stored only on DB0.</li> <li>The instance cannot be connected<br/>and remains read-only for 5 to 30<br/>minutes.</li> </ol> |
| Memcached | From single-<br>node to master/<br>standby  | Services are interrupted for several seconds and remain read-only for about 1 minute.  |

| Version           | Supported Type<br>Change  | Precautions  |
|-------------------|---|--|
| Redis 4.0/5.0/6.0 | From master/<br>standby or read/<br>write splitting to<br>Proxy Cluster                   | 1. Before changing the instance type to<br>Proxy Cluster, evaluate the impact on<br>services. For details, see What Are<br>the Constraints on Implementing   |
|                   | Proxy Cluster<br>From Proxy<br>Cluster to<br>master/standby<br>or read/write<br>splitting | <ul> <li>the Constraints on Implementing<br/>Multi-DB on a Proxy Cluster<br/>Instance? and Command<br/>Restrictions.</li> <li>Memory usage must be less than<br/>70% of the maximum memory of the<br/>new flavor.</li> <li>Some keys may be evicted if the<br/>current memory usage exceeds 90%<br/>of the total.</li> <li>After the change, create alarm rules<br/>again for the instance.</li> <li>For instances that are currently<br/>master/standby, ensure that their<br/>read-only IP address or domain name<br/>is not used by your application.</li> <li>If your application cannot reconnect<br/>to Redis or handle exceptions, you<br/>may need to restart the application<br/>after the change.</li> <li>Modify instance specifications during<br/>off-peak hours. An instance is</li> </ul> |
|                   |   | temporarily interrupted and remains<br>read-only for about 1 minute during<br>the specification change.  |

| Version           | Supported Type<br>Change   | Precautions  |
|-------------------|--|--|
| Redis 4.0/5.0/6.0 | From master/<br>standby to read/<br>write splitting<br>NOTE<br>Currently, a read/<br>write splitting<br>instance cannot<br>be directly<br>changed to a<br>master/standby<br>one. | <ol> <li>The instance memory must be<br/>greater than or equal to 4 GB, and<br/>will remain the same after the<br/>change.</li> <li>Some keys may be evicted if the<br/>current memory usage exceeds 90%<br/>of the total.</li> <li>After the change, create alarm rules<br/>again for the instance.</li> <li>Ensure that read-only IP addresses or<br/>domain names are not directly<br/>referred in the applications using the<br/>master/standby instance.</li> <li>If your application cannot reconnect<br/>to Redis or handle exceptions, you<br/>may need to restart the application<br/>after the change.</li> <li>Services may temporarily stutter<br/>during the change. Perform the<br/>change during off-peak hours.</li> <li>Unavailable for master/standby<br/>instances with ACL users.</li> <li>Unavailable for master/standby DCS<br/>Redis 6.0 instances with SSL enabled.</li> </ol> |

Any instance type changes not listed in the preceding table are not supported. To modify specifications while changing the instance type, create an instance, migrate data, and switch IPs. For details, see **Online Migration Between Instances**.

For details about the commands supported by different types of instances, see **Command Compatibility**.

#### Scaling

#### • Scaling options

| Table 8-2 | Scaling | options  | supported  | by diffe             | rent instances    |
|-----------|---------|----------|------------|----------------------|-------------------|
|           | ocaring | operoris | Jupporteur | <i>o</i> , <i>ac</i> | i chie mistariees |

| Cache<br>Engine | Single-<br>Node    | Master/<br>Standby  | Redis<br>Cluster | Proxy<br>Cluster | Read/<br>Write<br>Splitting |
|-----------------|--------------------|---------------------|------------------|------------------|-----------------------------|
| Redis 3.0       | Scaling<br>up/down | Scaling up/<br>down | -                | Scaling out      | -                           |

| Cache<br>Engine                           | Single-<br>Node    | Master/<br>Standby                                       | Redis<br>Cluster   | Proxy<br>Cluster               | Read/<br>Write<br>Splitting                                |
|---|--------------------|--|--|--------------------------------|--|
| Redis 4.0                                 | Scaling<br>up/down | Scaling up/<br>down and<br>replica<br>quantity<br>change | Scaling<br>up/down,<br>out/in,<br>and<br>replica<br>quantity<br>change | Scaling up/<br>down,<br>out/in | Scaling<br>up/down<br>and<br>replica<br>quantity<br>change |
| Redis 5.0                                 | Scaling<br>up/down | Scaling up/<br>down and<br>replica<br>quantity<br>change | Scaling<br>up/down,<br>out/in,<br>and<br>replica<br>quantity<br>change | Scaling up/<br>down,<br>out/in | Scaling<br>up/down<br>and<br>replica<br>quantity<br>change |
| Redis 6.0<br>basic<br>edition             | Scaling<br>up/down | Scaling up/<br>down and<br>replica<br>quantity<br>change | Scaling<br>up/down,<br>out/in,<br>and<br>replica<br>quantity<br>change | Scaling up/<br>down,<br>out/in | Scaling<br>up/down<br>and<br>replica<br>quantity<br>change |
| Redis 6.0<br>professio<br>nal<br>editions | -                  | Scaling up/<br>down                                      | -  | -                              | -  |
| Redis 7.0                                 | Scaling<br>up/down | Scaling up/<br>down and<br>replica<br>quantity<br>change | Scaling<br>up/down,<br>out/in,<br>and<br>replica<br>quantity<br>change | -                              | -  |
| Memcach<br>ed                             | Scaling<br>up/down | Scaling up/<br>down                                      | -  | -                              | -  |

#### 

- If the reserved memory of a DCS Redis 3.0 or Memcached instance is insufficient, the modification may fail when the memory is used up. For details, see **Reserved Memory**.
- Change the replica quantity and capacity separately.
- Only one replica can be deleted per operation.
- Impact of scaling

|--|

| Instance<br>Type                                  | Scali<br>ng<br>Type            | Impact   |
|---|--------------------------------|--|
| Single-<br>node ,<br>read/<br>write<br>splitting, | Scali<br>ng<br>up/<br>dow<br>n | • During scaling up, a basic edition DCS Redis 4.0 or later instance will be disconnected for several seconds and remain read-only for about 1 minute. During scaling down, connections will not be interrupted.   |
| and<br>master/<br>standby                         |                                | • A DCS Redis 3.0 instance will be disconnected for several seconds and remain read-only for 5 to 30 minutes.  |
|   |                                | • A DCS Redis professional edition instance will be disconnected for several seconds and remain read-only for about 1 minute.  |
|   |                                | • For scaling up, only the memory of the instance is expanded. The CPU processing capability is not improved.  |
|   | •                              | • Single-node DCS instances do not support data persistence. Scaling may compromise data reliability. After scaling, check whether the data is complete and import data if required. If there is important data, use a migration tool to migrate the data to other instances for backup. |
|   |                                | • For master/standby and read/write splitting instances, backup records created before scale-down cannot be used after scale-down. If necessary, download the backup file in advance or back up the data again after scale-down.   |

| Instance<br>Type                | Scali<br>ng<br>Type | Impact  |
|---------------------------------|---------------------|---|
| Proxy                           | Scali               | Scaling out by adding shards:   |
| Cluster<br>and Redis<br>Cluster | ng<br>up/<br>dow    | <ul> <li>Scaling out does not interrupt connections but<br/>will occupy CPU resources, decreasing<br/>performance by up to 20%.</li> </ul>  |
|                                 | n                   | <ul> <li>If the shard quantity increases, new Redis Server<br/>nodes are added, and data is automatically<br/>balanced to the new nodes, increasing the access<br/>latency.</li> </ul>  |
|                                 |                     | • Scaling in by reducing shards:  |
|                                 |                     | <ul> <li>If the shard quantity decreases, nodes will be<br/>deleted. Before scaling in a Redis Cluster<br/>instance, ensure that the deleted nodes are not<br/>directly referenced in your application, to<br/>prevent service access exceptions.</li> </ul>  |
|                                 |                     | <ul> <li>Nodes will be deleted, and connections will be<br/>interrupted. If your application cannot reconnect<br/>to Redis or handle exceptions, you may need to<br/>restart the application after scaling.</li> </ul>  |
|                                 |                     | • Scaling up by increasing the size per shard:  |
|                                 |                     | <ul> <li>Insufficient memory of the node's VM will<br/>cause the node to migrate. Service connections<br/>may stutter and the instance may become<br/>read-only during the migration.</li> </ul>  |
|                                 |                     | <ul> <li>Increasing the node capacity when the VM<br/>memory is sufficient does not affect services.</li> </ul>   |
|                                 |                     | <b>NOTE</b><br>Cluster DCS Redis 3.0 instances cannot be vertically scaled.   |
|                                 |                     | • Scaling down by reducing the shard size without changing the shard quantity has no impact.  |
|                                 |                     | • To scale down an instance, ensure that the used memory of each node is less than 70% of the maximum memory per node of the new flavor.  |
|                                 |                     | • The flavor changing operation may involve data migration, and the latency may increase. For a Redis Cluster instance, ensure that the client can process the MOVED and ASK commands. Otherwise, the request will fail.  |
|                                 |                     | <ul> <li>If the memory becomes full during scaling due to a<br/>large amount of data being written, scaling will fail.</li> </ul>   |
|                                 |                     | • Before scaling, <b>check for big keys through Cache</b><br><b>Analysis</b> . Redis has a limit on key migration. If the<br>instance has any single key greater than 512 MB,<br>scaling will fail when big key migration between<br>nodes times out. The bigger the key, the more likely<br>the migration will fail. |

| Instance<br>Type   | Scali<br>ng<br>Type  | Impact  |
|--|--|---|
|  |  | • Before scaling a Redis Cluster instance, ensure<br>that automated cluster topology refresh is<br>enabled. If it is disabled, you will need to restart the<br>client after scaling. For details about how to enable<br>automated refresh if you use Lettuce, see an<br>example of using Lettuce to connect to a Redis<br>Cluster instance.                         |
|  |  | <ul> <li>Backup records created before scaling cannot be<br/>used. If necessary, download the backup file in<br/>advance or back up the data again after scaling.</li> </ul>  |
| Master/<br>Standby,<br>read/<br>write<br>splitting,<br>and Redis<br>Cluster<br>instances | Scali<br>ng<br>out/i<br>n<br>(repli<br>ca<br>quan<br>tity<br>chan<br>ge) | • Before adding or removing replicas for a Redis<br>Cluster instance, ensure that automated cluster<br>topology refresh is enabled. If it is disabled, you will<br>need to restart the client after scaling. For details<br>about how to enable automated refresh if you use<br>Lettuce, see an example of using Lettuce to<br>connect to a Redis Cluster instance. |
|  |  | • Deleting replicas interrupts connections. If your application cannot reconnect to Redis or handle exceptions, you may need to restart the application after scaling. Adding replicas does not interrupt connections.  |
|  |  | <ul> <li>If the number of replicas is already the minimum<br/>supported by the instance, you can no longer delete<br/>replicas.</li> </ul>  |

#### **Changing an Instance**

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

- **Step 2** Click O in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Choose **More** > **Modify Specifications** of the **Operation** column in the row containing the DCS instance.
- **Step 5** On the **Modify Specifications** page, select the desired specification.

To expand the capacity of a single shard of a cluster instance, see **Can I Expand a Single Shard of a Cluster Instance?** 

**Step 6** Set **Apply Change** to **Now** or **During maintenance**.

Select **During maintenance** if the modification interrupts connections.

| Change  | When Connections Are Interrupted   |
|---|--|
| Scaling up a single-node<br>or master/standby<br>instance     | Memory is increased from a size smaller than 8 GB to 8 GB or larger.                               |
| Scaling down a Proxy<br>Cluster and Redis Cluster<br>instance | The number of shards is decreased.   |
| Changing the instance type                                    | The instance type is changed between master/<br>standby or read/write splitting and Proxy Cluster. |
| Deleting replicas   | Replicas are deleted from a master/standby, Redis<br>Cluster, or read/write splitting instance.    |

**Table 8-4** Scenarios where specification modification interrupts connections

#### **NOTE**

- If the modification does not interrupt connections, it will be applied immediately even if you select **During maintenance**.
- The modification cannot be withdrawn once submitted. To reschedule a modification, you can change the maintenance window. The maintenance window can be changed up to three times.
- Modifications on DCS Redis 3.0 and Memcached instances can only be applied immediately.
- If you apply the change during maintenance, the change starts at any time within the maintenance window, rather than at the start time of the window.
- If a large amount of data needs to be migrated when you scale down a cluster instance, the operation may not be completed within the maintenance window.
- Step 7 Click Next. Confirm the change details and view the risk check results.

If any risk is found in the check, the instance may fail to be modified. For details, see **Table 8-5**.

| Table 8-5 Risk check iter | ns |
|---------------------------|----|
|---------------------------|----|

| Check Item  | Reason for Check  | Solution   |
|---|---|--|
| <ul> <li>Non-standard<br/>configuration check</li> <li>NOTE         <ul> <li>Currently, non-standard<br/>configuration check is<br/>available only in some<br/>regions, such as CN<br/>North-Beijing4, CN<br/>East-Shanghai1, and<br/>CN East-Shanghai2.</li> <li>Check whether the<br/>following items meet<br/>standards:</li></ul></li></ul> | If your instance has non-<br>standard configurations,<br>the console displays a<br>message indicating that<br>they will be converted to<br>standard during the<br>change.<br>You can retain non-<br>standard bandwidth or<br>proxy quantity<br>configuration only.  | <ul> <li>If your instance does<br/>not have non-<br/>standard<br/>configurations, the<br/>check result is normal<br/>and no action is<br/>required.</li> <li>If the instance has<br/>non-standard<br/>configurations,<br/>determine whether to<br/>proceed with the<br/>change or whether to<br/>retain the non-<br/>standard bandwidth<br/>and proxy quantity<br/>configuration.</li> </ul> |
| Node status   | Abnormal instance nodes<br>cause instance<br>modification failures.   | If this case, contact customer service.  |
| Dataset memory<br>distribution check<br><b>NOTE</b><br>This check item applies<br>only to Proxy Cluster and<br>Redis Cluster instances.   | Specification<br>modification of a cluster<br>instance involves data<br>migration between<br>nodes. If an instance has<br>any key bigger than 512<br>MB, the modification will<br>fail when big key<br>migration between<br>nodes times out.<br>If the instance dataset<br>memory is unevenly<br>distributed among nodes<br>and the difference is<br>greater than 512 MB, the<br>instance has a big key<br>and the change may fail. | Handle big keys before<br>proceeding with the<br>change.   |

| Check Item   | Reason for Check  | Solution   |  |
|--|---|--|--|
| Memory usage check   | If the memory usage of<br>a node is greater than<br>90%, keys may be<br>evicted or the change<br>may fail.  | If the memory usage is<br>too high, optimize the<br>memory by optimizing<br>big keys, scanning for<br>expired keys, or deleting<br>some keys.  |  |
| Network input traffic<br>check<br><b>NOTE</b><br>This check item applies<br>only to single-node, read/<br>write splitting, and master/<br>standby instances. | The change may fail if<br>the network input traffic<br>is too heavy and the<br>write buffer overflows.  | Perform the change<br>during off-peak hours.   |  |
| CPU usage check  | If the node CPU usage<br>within 5 minutes is<br>greater than 90%, the<br>change may fail.   | Perform the change<br>during off-peak hours.<br>Troubleshooting High<br>CPU Usage of a DCS<br>Redis Instance   |  |
| Resource capacity<br>NOTE<br>This item should be<br>checked only when scaling<br>up cluster instances.   | To scale up a cluster<br>instance, if the VM<br>resource capacity is<br>insufficient, the node<br>needs to be migrated<br>during the change.<br>Service connections<br>become intermittent or<br>read-only during the<br>migration. | If the resource capacity<br>check poses risks, ensure<br>that your application can<br>reconnect to Redis or<br>handle exceptions, you<br>may need to restart the<br>application after the<br>change. |  |

#### 

- If the check results are normal, no risks are found in the check.
- If the check fails, the possible causes are as follows:
  - The master node of the instance fails to be connected. In this case, check the instance status.
  - The system is abnormal. In this case, click **Check Again** later.
- Click **Stop Check** to stop the check. Click **Check Again** to restart the check.
- If you want to proceed with the change despite risks found in the check or after clicking **Stop Check**, select **I understand the risks.**
- **Step 8** After the risk check is complete, click **Next**. After the modification is submitted, you can go to the **Background Tasks** page to view the modification status.

Click the task name on the **Background Tasks** page to view task details. After an instance is successfully modified, it changes to the **Running** state.

|                                  |   |  |   |   | Task Details  |  |  |
|----------------------------------|---|--|---|---|---|--|--|
| 1. A back<br>shown<br>2. After n | ground task starts when y<br>here.<br>nodifying the specification | rou perform an operatio<br>is, configure alarm rules | n such as modifying the spec<br>on this page. | fications, enabling/disabling public acce | Progress  | 100 %  |  |
|                                  |   |  |   |   | Steps   |  |  |
| No                               | Task Name   | Usemame  | Chature                                       | Start Time                                | 1. Pre-check Mar 22, 2023 09:44:26 ~                  | Mar 22, 2023 09:44:29                            |  |
| NO.                              | Task Name   | Username   | Status  | start lime                                | 2. Create New Nodes Mar 22, 2023                      | 09:44:38 ~ Mar 22, 2023 09:44:53 💙               |  |
| 1                                | Modify specification  |  | Successful                                    | Mar 22, 2023 09:44:25 GMT+08              | 3. Migrate Data Mar 22, 2023 09:44:                   | 59~Mar 22, 2023 09:45:02                         |  |
|                                  |   |  |   |   | Ø Migrate full data                                   | 📀 Migrate incremental data                       | Set read-only                          |
|                                  |   |  |   |   | Mar 22, 2023 09:44:59 ~<br>Mar 22, 2023 09:44:59      | Mar 22, 2023 09:44:59 ~<br>Mar 22, 2023 09:44:59 | Mar 22, 2023 09:4<br>Mar 22, 2023 09:4 |
|                                  |   |  |   |   | <ul> <li>4. Configure New Nodes Mar 22, 20</li> </ul> | 023 09:45:02 ~ Mar 22, 2023 09:45:05 💙           |  |
|                                  |   |  |   |   | 5 Delete Original Nodes Mar 22.2                      | 023 09:45:09 ~ Mar 22, 2023 09:45:24             |  |

0 4 V/ · . . . . . .

#### **NOTE**

- If the specification modification of a single-node DCS instance fails, the instance is temporarily unavailable. The specification remains unchanged. Some management operations (such as parameter configuration and specification modification) are temporarily not supported. After the specification modification is completed in the backend, the instance changes to the new specification and becomes available for use again.
- If the specification modification of a master/standby or cluster DCS instance fails, the instance still uses its original specifications. Some management operations (such as parameter configuration, backup, restoration, and specification modification) are temporarily not supported. Remember not to read or write more data than allowed by the original specifications; otherwise, data loss may occur.
- After the specification modification is successful, the new specification of the instance takes effect.
- Specification modification of a single-node, master/standby, or read/write splitting DCS instance takes 5 to 30 minutes to complete, while that of a cluster DCS instance takes a longer time.

----End

#### 8.2 Adjusting DCS Instance Bandwidth

Generally, Redis instances save and obtain data in the data layer closer to application services, which consumes the network bandwidth. Rate limits may occur when the instance bandwidth is insufficient, causing increased service latency or client connection exceptions. Currently, the Redis instance bandwidth can be adjusted on the console for DCS Redis 4.0 and later instances.

#### Notes and Constraints

- This function is unavailable for professional edition DCS Redis instances. •
- This function is available only for instances in the **Running** state. •
- The adjustment range of bandwidth per shard is from the shard's assured • (default) bandwidth to its maximum bandwidth. Generally, the maximum bandwidth per shard is 2,048 Mbit/s when the physical machine of the instance node has sufficient resources.
- Set the target bandwidth to a multiple of 8. Otherwise, the value will be automatically rounded down to a multiple of 8 after the order is submitted.
- Changed bandwidth is pay-per-use only, and charged by hour. Check your fee changes.

#### Procedure

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

- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- Step 3 In the navigation pane, choose Cache Manager.
- **Step 4** On the **Cache Manager** page, click the name of the DCS instance you want to configure.
- **Step 5** In the **Instance Details** area of the DCS instance, click **Adjust Bandwidth** next to **Bandwidth**.

| Figure 8-2 / | djusting | bandwidth |
|--------------|----------|-----------|
|--------------|----------|-----------|

| Instance Details           |                            |
|----------------------------|----------------------------|
| Name                       | dcs 2                      |
| Status                     | Running     ■              |
| ID                         | aecf O                     |
| Cache Engine               | Basic   Redis 4.0          |
| Instance Type              | Master/Standby             |
| Cache Size                 | 0.125 GB   Replicas: 2     |
| Bandwidth                  | 40 Mbit/s Adjust Bandwidth |
| Used/Available Memory (MB) | 2/128 (1.56%) ⑦            |

Step 6 On the Adjust Bandwidth page, set bandwidth parameters.

By default, the bandwidth can be manually adjusted. Enabling the auto scaling whitelist supports both **Manual** and **Auto scaling**. If the adjustment cannot be specified on the console, contact customer service.

#### **NOTE**

- Set the target bandwidth to a multiple of 8. Otherwise, a value rounded down to a multiple of 8 will be automatically used after the order is submitted. For example, if you set the bandwidth to 801, 800 will be used instead.
- Price shown on the change page is only the fee of the additional bandwidth.
- The additional bandwidth is pay-per-use, and charged by hour.
- You can adjust the bandwidth whenever as required. If you perform multiple bandwidth changes in a billing period (one hour), you will be billed according to highest bandwidth in the period. For example, if you have changed the bandwidth of a DCS Redis instance from 256 Mbit/s (default) to 2,048 Mbit/s, and changed the bandwidth again to 512 Mbit/s in the same billing period, you will be billed at the price of 2,048 Mbit/s bandwidth.
- Manual
  - a. Set target bandwidth. For cluster instances with multiple shards, specify the target for each of them, or select the desired shards and click **Adjust Bandwidth**.

#### Figure 8-3 Manually setting target bandwidth

| Adjust Bandwidth   |  |  |  |  |  |
|--|--|--|--|--|--|
| <ol> <li>Set the target bandwidth to a multiple of 8.Otherwise,</li> <li>The instance bandwidth is pay-per-use only. Check your</li> </ol> | the value will be automatically rounded dowr<br>fee changes. | to a multiple of 8 after the order is submitted. |  |  |  |
| Mode Manual Auto scaling Set different bandwidths for different shards, or set the same bandwidth Adjust Bandwidth                         | width for multiple shards in batches.                        |  |  |  |  |
| 🕑 Shard Name 🔶   | Assured (Mbit/s) 🔶 Curre                                     | nt (Mbit/s) \ominus 🛛 Target (Mbit/s) \ominus    |  |  |  |
| group-0<br>88b0a321-8b23-4325-aee5-2131485cae4f  | 768  | 768 - 768 +                                      |  |  |  |
| group-1<br>867c3671-f096-4cf1-a3ee-e0a7b3efc602  | 768  | 768 - 768 +                                      |  |  |  |
| group-2<br>20766148-509e-4118-a36c-7aa3fe5f4381  | 768  | 768 - 768 +                                      |  |  |  |

- b. Confirm the bandwidth and fees, check Authorization, and click Submit.
   When the bandwidth adjustment task is in the Successful state, the new bandwidth is used.
- Auto scaling
  - a. Enable **Auto Bandwidth Increase** and set the policies as required, as shown in **Table 8-6**.

Bandwidth increases automatically (up to 2,048 Mbit/s per shard) based on scaling policies. Automatic scaling overrides manual adjustments.

#### Figure 8-4 Setting auto bandwidth increase policies

| Mode                |                                     |  |
|---------------------|-------------------------------------|--|
| Manual              | Auto scaling                        |  |
| Bandwidth increases | automatically (up to 2,048 Mbit/s p | per shard) based on scaling policies. Automatic scaling overrides manual adjustmen |
| Auto Bandwidth Incr | ise                                 |  |
| Enabled             |                                     |  |
| Once burst bandwid  | usage reaches the specified valu    | e for a monitoring period, the bandwidth automatically increases.                  |
| Burst Bandwidth Us  | e≥ ⑦                                |  |
| 50                  | ~ %                                 |  |
| Burst bandwidth usa | e = Max (Input flow, Output flow)/S | Shard bandwidth  |
| Monitoring Period ( |                                     |  |
| 1                   | ✓ min                               |  |
| Silence (?)         |                                     |  |
| 0                   | seconds                             |  |
| Range: 0 to 86400   |                                     |  |
| Authorization       |                                     |  |
| 🗸 I am aware of th  | expense changes caused by the       | change and agree to implement it.  |
|                     |                                     |  |

| Policy                     | Description  |  |
|----------------------------|--|--|
| Burst Bandwidth<br>Usage ≥ | Burst bandwidth usage threshold for bandwidth increases.   |  |
|                            | Calculation:   |  |
|                            | Burst bandwidth usage = Burst used bandwidth/<br>Shard bandwidth. The value of "Burst used<br>bandwidth" uses the larger metric between<br><b>Output Flow</b> and <b>Input Flow</b> .  |  |
|                            | Target:  |  |
|                            | When the burst bandwidth usage of an instance<br>shard reaches the threshold, the shard bandwidth<br>is automatically scaled up. As a result, the burst<br>bandwidth usage is reduced to (its threshold<br>minus 10%).   |  |
|                            | For example, if this threshold is set to 70%, the<br>shard bandwidth will be automatically increased<br>when the burst bandwidth usage reaches 70%.<br>As a result, the burst bandwidth usage decreases<br>to 60%. Therefore, the shard bandwidth after the<br>scaling equals the burst bandwidth usage divided<br>by 60%. |  |
| Monitoring Period          | Monitoring period of bandwidth increases, in minutes. Default: 1   |  |
|                            | For example, if the monitoring period is set to 1<br>minute, the bandwidth data is monitored within<br>1 minute.   |  |
| Silence                    | Interval between increases, in seconds. Default: 0<br>The silence time avoids consecutive automatic<br>bandwidth increases.  |  |

|  | Table 8-6 | Setting | auto | bandwidth | increase | policies |
|--|-----------|---------|------|-----------|----------|----------|
|--|-----------|---------|------|-----------|----------|----------|

b. Confirm the bandwidth parameters, check **Authorization**, and click **Submit**.

----End

#### Checking Assured Bandwidth and Adjusted Bandwidth

On the page for manually adjusting the bandwidth, you can view **Assured** and **Current** bandwidth of each shard. **Current** displays the latest bandwidth.

#### Figure 8-5 Viewing bandwidth

| Adjust Bandwidth                                |                  |                  |                 |
|---|------------------|------------------|-----------------|
| Shard Name 🔶                                    | Assured (Mbit/s) | Current (Mbit/s) | Target (Mbit/s) |
| group-0<br>fe4912d3-de7e-4e8c-b265-e3264b7ba734 | 768              | 800              | - 800 +         |
| group-1<br>dc44b099-1954-4b37-b8ce-e8cec8e9dd7a | 768              | 800              | - 800 +         |
| group-2<br>ec99dcd9-e779-4c68-b989-d66862e4e832 | 768              | 800              | - 800 +         |

The relationship between the instance bandwidth and the bandwidth of a single shard is as follows:

- Bandwidth of single-node or master/standby instances = Bandwidth per shard
- Bandwidth of read/write splitting instances = Bandwidth per shard × Replica quantity
- Bandwidth of cluster instances = Bandwidth per shard × Shard quantity, or the total bandwidth of all shards if the bandwidth per shard varies

For example, **Figure 8-5** shows a cluster instance with three shards. The adjusted bandwidth of each shard is 800 Mbit/s, and the total bandwidth of the cluster instance is 2,400 Mbit/s.

#### 8.3 Changing Cluster DCS Instances to be Across AZs

To implement disaster recovery, cluster instances (whose master and standby nodes are) in a single AZ can be deployed across AZs by migrating the standby nodes to other AZs.

#### Notes and Constraints

- Available only for single-AZ cluster instances with two or more replicas.
- When you enable multi-AZ for a Proxy Cluster instance:
  - Service running may fluctuate during the change. Perform this operation during off-peak hours.
  - If your application cannot reconnect or handle exceptions, try restarting the application after the change.
- When you enable multi-AZ for a Redis Cluster instance:
  - Changing AZs will not interrupt services or the master node, but will slightly affect performance. Perform this operation during off-peak hours.
  - Changing AZs interrupts connections to some replicas. Ensure your application can automatically recover from exceptions and reconnect to Redis.

#### Procedure

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

**Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.

- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** On the **Cache Manager** page, click a DCS instance.
- Step 5 In the Network area of the DCS instance, click Enable Multi-AZ next to AZ.

Figure 8-6 Enabling multi-AZ

| Instance Details           |  | Connection ③           |                               |
|----------------------------|--|------------------------|-------------------------------|
| Name                       | dcs-036t 🖉   | Password Protected     | No                            |
| Status                     | S Running  | Connection Address     | redis                         |
| ID                         | (9a)   | IP Address             | 192.168.                      |
| Cache Engine               | Redis 5.0  | Public Access          | View Documentation            |
| Instance Type              | Proxy Cluster  | Client IP Pass-through | off∉ ⑦                        |
| Cache Size                 | 4 GB   Replicas: 2   Shards: 3   Shard Size: 1.33 GB |                        |                               |
| Bandwidth                  | 2,304 Mbit/s Increase Bandwidth Temporarily          |                        |                               |
| Used/Available Memory (MB) | 9/4,096 (0.22%) ③                                    | Network                |                               |
| CPU                        | x86  | AZ                     | AZ1 Enable Multi-AZ           |
| Enterprise Project         | default <u>2</u>                                     | VPC                    | vpc-05-qianyi-test 🕜          |
| Maintenance                | 02:00 03:00 @ GMT+08:00                              | Subnet                 | subnet-383b(192.168.100.0/24) |

Step 6 In the displayed Change AZs dialog box, specify Standby AZ.

Figure 8-7 Selecting a standby AZ

| ID           | 89d       |   |
|--------------|-----------|---|
| Name         | dcs       |   |
| Primary AZ   |           |   |
| Standby AZ   | Select    | ~ |
| Apply Change | Now       | ~ |
|              | OK Cancel |   |

- Step 7 Set Apply Change to Now or During maintenance.
- Step 8 Click OK.

When complete, the task changes to the **Successful** state.

----End

### 8.4 Upgrading Minor or Proxy Versions of a DCS Instance

DCS optimizes functions and fixes vulnerabilities in minor and proxy upgrades. This section describes how to upgrade the minor or proxy version of an instance. Upgrading the minor or proxy version does not affect the instance connection addresses, password, whitelist, or monitoring and alarms.

#### **NOTE**

Currently, this function is whitelist-based. To enable it, contact customer service.

#### Notes and Constraints

- This function is available only for DCS Redis 4.0 or later basic edition instances.
- Only Proxy Cluster and read/write splitting instances involve proxy versions.
- If the minor or proxy version of the instance is already the latest, you cannot upgrade it.
- To upgrade the minor version of a Redis Cluster instance, ensure that the client can properly process the **MOVED** and **ASK** commands. Otherwise, requests will fail.

#### Impacts

- Perform instance upgrades during off-peak hours. Otherwise, upgrades may fail when the instance memory or CPU usage exceeds 90% or write traffic bursts. In such cases, try again during off-peak hours.
- The instance's minor version is upgraded by migrating nodes. During the migration, latency will increase. A migrating shard will become read-only for 1 minute and intermittently disconnected. Ensure that the client can reconnect and handle exceptions.
- The upgrading instance will be intermittently disconnected. Ensure that the client can reconnect and handle exceptions. Perform the upgrade during off-peak hours.

#### **Upgrading Minor or Proxy Versions of a DCS Instance**

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

- **Step 2** Click <sup>(2)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click a DCS instance.
- **Step 5** On the **Basic Information** page, you can view or upgrade the minor or proxy versions of the instance.

| 3 13             | 5                    | 1 2  |
|------------------|----------------------|------|
| Instance Details |                      |      |
| Name             | dcs-                 | 2    |
| Status           | 😔 Running            |      |
| ID               | 2ec1833e             | )c 🗍 |
| Cache Engine     | Basic   Redis 5.0    |      |
| Minor Version    | 5.0.12.0 Upgrade     |      |
| Proxy Version    | 5.0.14.0 Upgrade     |      |
| Instance Type    | Read/Write splitting |      |
| Cache Size       | 2 GB   Replicas: 3   |      |

Figure 8-8 Upgrading an instance's minor or proxy version

- Upgrading a minor version
  - a. Click Upgrade next to Minor Version.

Also, to upgrade the proxy version, enable **Upgrade Proxy Version** in the displayed window.

- b. Click **OK**. The upgrade is complete when the upgrade task is in the **Successful** state.
- Upgrading a proxy version
  - a. Click Upgrade next to Proxy Version.
    - Also, to upgrade the minor version, enable **Upgrade Minor Version** in the displayed window.
  - b. Click **OK**. The upgrade is complete when the upgrade task is in the **Successful** state.

----End

#### **Related Documents**

See the DCS Redis Release History.

## **9** Managing Lifecycle of an Instance

#### 9.1 Restarting a DCS Instance

To recover an instance in cases such as high memory fragmentation ratio or fault occurrence, try restarting the instance on the DCS console. DCS instances can be restarted in batches.

#### **Notes and Constraints**

- The DCS instances must be in the **Running** or **Faulty** state.
- For single-node instances or master/standby, cluster, and read/write splitting ones with AOF persistence disabled (parameter **appendonly** set to **no**), the instance data will be cleared after an instance restart. Exercise caution.
- While a DCS instance is restarting, it cannot be read or written.
- An attempt to restart a DCS instance while it is being backed up cancels the backup task, or may result in a failure.
- Restarting a DCS instance will disconnect the original client. You are advised to configure automatic reconnection in your application.

#### Procedure

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click <sup>(2)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Select the names of one or more desired DCS instances and click **Restart** above the list.

To restart one instance, you can also locate the row containing the desired instance, and click **Restart** in the **Operation** column.

**Step 5** In the displayed dialog box, click **Yes**. After DCS instances are restarted, their status changes to **Running**.

The time required for restarting a DCS instance depends on the cache size of the instance. It make take **10s to 30 minutes**.

#### D NOTE

By default, only the instance process is restarted. If you select **Force restart** for a DCS Redis 3.0 or Memcached instance, its VM will be restarted. **Force restart** is not supported by DCS Redis 4.0 or later instances.

----End

#### 9.2 Starting or Stopping a DCS Instance

Redis 4.0 and later instances support instance stop. When an instance is stopped, data reading or writing is stopped so that the instance cannot be modified, configured, backed up, or migrated. You can neither change the password nor analyze the cache.

A Redis instance is in **Running** state by default, or is in **Stopped** state after you stop it. **Stopping an instance does not affect its billing.** 

#### Procedure

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Stop or start DCS instances.
  - Stopping instances
    - a. Locate the desired instance and choose More > Stop in the Operation column. You can also select desired instances and choose More > Stop above.
    - b. A dialog box is displayed. Click **Yes**. When the instance is in **Stopped** state, the instance is stopped.
  - Starting instances
    - a. To restart an instance, click Start in the Operation column of the desired instance, or select desired instances on the left and choose More > Start above.
    - b. A dialog box is displayed. Click **Yes**. The instance is started when it is in the **Running** state.

----End

#### 9.3 Deleting a DCS Instance

On the DCS console, you can delete one or multiple DCS instances at a time. You can also delete all instance creation tasks that have failed to run.

#### **Notes and Constraints**

- The DCS instances exist, and must be in the **Running**, **Faulty**, or **Stopped** state.
- Deleting DCS instances removes their data and backups permanently. To retain the backups, download and save them first.
- If the instance is in cluster mode, all cluster nodes will be deleted.
- Instances billed on a yearly/monthly basis cannot be deleted.

#### Procedure

#### **Deleting Successfully Created Instances**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click <sup>(Q)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** On the **Cache Manager** page, select one or more DCS instances you want to delete.

DCS instances in the **Creating**, **Starting**, **Stopping**, or **Restarting** state cannot be deleted.

- Step 5 Choose More > Delete above the instance list.
- **Step 6** Enter **DELETE** and click **Yes** to delete the DCS instance.

It takes 1 to 30 minutes to delete DCS instances.

#### **NOTE**

To delete a single instance, choose **More** > **Delete** in **Operation** column in the row containing the instance.

#### ----End

#### **Deleting Instances That Failed to Be Created**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click <sup>(2)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.

If there are DCS instances that have failed to be created, **Instance Creation Failures** and the number of instances that fail to be created is displayed above the instance list.

**Step 4** Click the icon or the number of failed tasks next to **Instance Creation Failures**.

The Instance Creation Failures dialog box is displayed.

- **Step 5** Delete failed instance creation tasks as required.
  - To delete all failed tasks, click **Delete All** above the task list.

• To delete a single failed task, click **Delete** in the row containing the task.

----End

#### 9.4 Clearing DCS Instance Data

To clear instance data, use the **FLUSHDB** or **FLUSHALL** commands on accessed instances, or the data clearance function on the DCS console. This section describes how to use the function to clear instance data with one click.

#### NOTICE

Clearing instance data may cause the service latency to increase sharply.

#### **Notes and Constraints**

- The instances must be of Redis 4.0 and later, and in the **Running** state.
- Clearing instance data cannot be undone and cleared data cannot be recovered. Exercise caution when performing this operation.

#### Procedure

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Select one or more DCS instances.
- **Step 5** Choose **More** > **Clear Data** above the instance list.
- **Step 6** In the displayed dialog box, click **Yes**.

----End

# **10** Diagnosing and Analyzing an Instance

#### 10.1 Querying Big Keys and Hot Keys in a DCS Redis Instance

Big keys and hot keys are common issues. This section describes the big key and hot key analysis function on the DCS console. This function monitors the key that occupies most space of a Redis instance, or that is most frequently accessed from the storage data.

- There are two types of big keys:
  - The key value occupies much storage space. If the size of a single String key exceeds 10 KB, or if the size of all elements of a key combined exceeds 50 MB, the key is defined as a big key.
  - The key contains many elements. If the number of elements in a key exceeds 5000, the key is defined as a big key.
- A hot key is most frequently accessed, or consumes significant resources. For example:
  - In a cluster instance, a shard processes 10,000 requests per second, among which 3000 are performed on the same key.
  - In a cluster instance, a shard uses a total of 100 Mbits/s inbound and outbound bandwidth, among which 80 Mbits/s is used by the HGETALL operation on a Hash key.

#### **Notes and Constraints**

Big key and hot key analysis consumes CPU. Perform big key and hot key analysis during off-peak hours to avoid 100% CPU usage.

#### Notes on big key analysis:

• During big key analysis, all keys will be traversed. The larger the number of keys, the longer the analysis takes.

- Perform big key analysis during off-peak hours and avoid automatic backup periods.
- For a master/standby, read/write splitting, or cluster instance, the big key analysis is performed on the standby node, so the impact on the instance is minor. For a single-node instance, the big key analysis is performed on the only node of the instance and will reduce the instance access performance.
- A maximum of 100 analysis records are retained for each instance. When this limit is reached, the oldest records will be deleted to make room for new records. You can also manually delete records you no longer need.

#### Notes on hot key analysis:

- Hot keys can be analyzed only for DCS Redis 4.0 and later instances.
- The **maxmemory-policy** parameter of the instance must be set to **allkeys-lfu** or **volatile-lfu**.
- During hot key analysis, all keys will be traversed. The larger the number of keys, the longer the analysis takes.
- Perform hot key analysis shortly after peak hours to ensure the accuracy of the analysis results.
- The hot key analysis is performed on the master node of each instance and will reduce the instance access performance.
- A maximum of 100 hot key analysis records are retained for each instance. When this limit is reached, the oldest records will be deleted to make room for new records. You can also manually delete records you no longer need.

#### Querying Big Keys

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

- **Step 2** Click Sin the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of a DCS Redis instance.
- **Step 5** Choose **Analysis and Diagnosis** > **Cache Analysis**.
- **Step 6** On the **Big Key Analysis** tab page, you can manually start a big key analysis or schedule a daily automatic analysis.
- **Step 7** After an analysis task completes, click **View** to view the analysis results of different data types.

You can also click **Download** or **Delete** in the **Operation** column to download or delete the analysis result.

**NOTE** 

The big key analysis result shows records of the top 100 (20 for Strings and 80 for Lists/ Sets/Zsets/Hashes) data size.

#### Figure 10-1 Viewing the results of big key analysis (for Strings)

Analysis Task Details

| Task ID    | 249d47bf-61f8-40df-9789-985645f390ef | Start Time | Jul 08, 2024 11:2 | 2:26 GMT+08:00 |          |
|------------|--------------------------------------|------------|-------------------|----------------|----------|
| Status     | Successful                           | End Time   | Jul 08, 2024 11:2 | 2:57 GMT+08:00 |          |
| Strings    | Lists/Sets/Zsets/Hashes              |            |                   |                |          |
| 20 records |                                      |            |                   |                | Q        |
| Key        |                                      | Ту         | rpe               | Bytes          | Database |
| normal-8   | 73                                   | St         | ring              | 3,128          | 0        |
| normal-2   | 294                                  | St         | ring              | 3,128          | 0        |
| normal-1   | 949                                  | St         | ring              | 3,128          | 0        |
| normal-1   | 616                                  | St         | ring              | 3,128          | 0        |
| normal-4   | 5                                    | St         | ring              | 3,128          | 0        |
| normal-1   | 30                                   | St         | ring              | 3,128          | 0        |
| normal-1   | 663                                  | St         | ring              | 3,128          | 0        |
| normal-1   | 347                                  | St         | ring              | 3,128          | 0        |
| normal-2   | 815                                  | St         | ring              | 3,128          | 0        |
| normal-7   | 96                                   | St         | ring              | 3,128          | 0        |
|            |                                      |            |                   |                |          |

Figure 10-2 Viewing the results of big key analysis (for Lists/Sets/Zsets/Hashes)

| Analysis Task Details                 |  |                    |               |          |
|---------------------------------------|--|--------------------|---------------|----------|
| Task ID 249d47bf-61f8-40df-9789-98564 | 15f390ef Start Time  | Jul 08, 2024 11:22 | :26 GMT+08:00 |          |
| Status 😔 Successful                   | Status 📀 Successful End Time Jul 08, 2024 11:22:57 GMT+08:00 |                    |               |          |
| Strings Lists/Sets/Zsets/Hashes       |  |                    |               |          |
| 80 records                            |  |                    |               | Q        |
| Кеу                                   | Туре   | Bytes              | Quantity      | Database |
| stream-858                            | Stream   | 88323              | 24            | 0        |
| stream-476                            | Stream   | 73670              | 15            | 0        |
| hash-858                              | Hash   | 68875              | 24            | 0        |
| hash-476                              | Hash   | 66688              | 15            | 0        |
| hash-1325                             | Hash   | 65694              | 12            | 0        |
| hash-2210                             | Hash   | 65694              | 12            | 0        |
| stream-322                            | Stream   | 64782              | 15            | 0        |
| stream-1325                           | Stream   | 63143              | 12            | 0        |
| stream-757                            | Stream   | 62456              | 18            | 0        |
| 858                                   | Zset   | 59417              | 24            | 0        |
|                                       |  |                    |               |          |

| Parameter | Description   |
|-----------|---|
| Key       | The key name in a big key analysis result.  |
| Туре      | Type of a key, which can be string, hash, list, set, or zset.   |
| Size      | The value size of a key, in Bytes.  |
| Quantity  | Number of elements in a key. This parameter is<br>displayed only for list, set, zset, or hash type. Unit:<br>counts |
| Database  | Database where the key is located.  |

| Table 10-1 Results o | f big key analysis |
|----------------------|--------------------|
|----------------------|--------------------|

----End

#### **Querying Hot Keys**

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

- **Step 2** Click Sin the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of a DCS Redis instance.
- **Step 5** Choose Analysis and Diagnosis > Cache Analysis.
- **Step 6** On the **Hot Key Analysis** tab page, you can manually start a hot key analysis or schedule a daily automatic analysis.

#### 

If the instance was created before July 2020, the default value of the **maxmemory-policy** parameter is **noeviction**. If the instance was created in or after July 2020, the default value of the **maxmemory-policy** parameter is **volatile-lru**. To perform hot key analysis, set this parameter to **allkeys-lfu** or **volatile-lfu** on the **Instance Configuration** > **Parameters** page. For details about **allkeys-lfu** and **volatile-lfu**, see **What Is the Default Data Eviction Policy**?

Step 7 After an analysis task completes, click View to view the analysis results.

You can also click **Download** or **Delete** in the **Operation** column to download or delete the analysis result.

#### **NOTE**

The hot key analysis result shows the most 100 frequently accessed keys within the specified period.

| analysis Task De  | tails   |                         |                      |          |
|---|---|-------------------------|----------------------|----------|
| ask ID 71200223-  | bada-4905-842f-9bfce3bea6ea   | Start Time Jul 08, 2024 | 4 11:21:45 GMT+08:00 |          |
| tatus 📀 Succe   | ssful   | End Time Jul 08, 2024   | 4 11:21:47 GMT+08:00 |          |
| 00 records  |   |                         |                      |          |
| Кеу   | Туре  | Size                    | FREQ (?)             | Database |
| (an HD  | Traper.   | 15,073 Bytes            | 8                    | 0        |
| (herea)   | in parts  | 1,998 Bytes             | 8                    | 0        |
| penetra   | (minor)   | 23,261 Bytes            | 8                    | 0        |
| (ineq)  | Traper.   | 28,178 Bytes            | 8                    | 0        |
| pendi .   | in parts  | 22,544 Bytes            | 8                    | 0        |
| part of the second  | (minute)  | 388 Bytes               | 7                    | 0        |
| per lli   | The party   | 3,480 Bytes             | 7                    | 0        |
| per se la companya de | Depen   | 2,867 Bytes             | 7                    | 0        |
| per fini  | inter-  | 8,673 Bytes             | 7                    | 0        |
| and the   | Traper-   | 11,244 Bytes            | 7                    | 0        |
|   | the second se | 10 100 0 1              | _                    |          |

Figure 10-3 Viewing the results of hot key analysis

| Table | 10-2 | Results | of | hot | key | analysis |
|-------|------|---------|----|-----|-----|----------|
|-------|------|---------|----|-----|-----|----------|

| Parameter | Description  |
|-----------|--|
| Кеу       | The key name in a hot key analysis result.   |
| Туре      | Type of a key, which can be string, hash, list, set, or zset.  |
| Size      | The value size of a key, in Bytes.   |
| FREQ      | Reflects the access frequency of a key within a specific period of time (usually 1 minute).  |
|           | <b>FREQ</b> is the logarithmic access frequency counter. The maximum value of <b>FREQ</b> is 255, which indicates 1 million access requests. After <b>FREQ</b> reaches <b>255</b> , it will no longer increment even if access requests continue to increase. <b>FREQ</b> will decrement by 1 for every minute during which the key is not accessed. |
| Shard     | Shard where the key is located.<br><b>NOTE</b><br>This parameter is displayed only for cluster instances.  |
| Database  | Database where a key is located.   |

----End

#### FAQs About Big Keys and Hot Keys

- Why Is the Capacity or Performance of a Shard of a Redis Cluster Instance Overloaded When That of the Instance Is Still Below the Bottleneck?
- What Is the Impact of a Hot Key?
- How Do I Avoid Big Keys and Hot Keys?
- How Do I Analyze the Hot Keys of a DCS Redis 3.0 Instance?
- How Do I Detect Big Keys and Hot Keys in Advance?

## 10.2 Scanning and Deleting Expired Keys in a DCS Redis Instance

There are two ways to delete a key in Redis.

- Use the **DEL** command to directly delete a key.
- Use commands such as **EXPIRE** to set a timeout on a key. After the timeout elapses, the key becomes inaccessible but is not deleted immediately because Redis is mostly single-threaded. Redis uses the following strategies to release the memory used by expired keys:
  - Lazy free deletion: The deletion strategy is controlled in the main I/O event loop. Before a read/write command is executed, a function is called to check whether the key to be accessed has expired. If it has expired, it will be deleted and a response will be returned indicating that the key does not exist. If the key has not expired, the command execution resumes.
  - Scheduled deletion: A time event function is executed at certain intervals. Each time the function is executed, a random collection of keys are checked, and expired keys are deleted. Instead of checking all keys each time, open-source Redis randomly checks 20 keys each time, 10 times per second by default. This avoids prolonging blocks on the Redis main thread, but the memory used by expired keys cannot be released quickly.

DCS integrates these strategies, and provides a common expired key query method to allow you to periodically release the memory used by expired keys. You can configure scheduled scans on the master nodes of your instances. The entire keyspace is traversed during the scans, triggering Redis to check whether the keys have expired and to remove expired keys if any.

#### **Notes and Constraints**

- Expired keys can be scanned only for DCS Redis 4.0 and later (enterprise edition excluded) instances.
- Released expired keys cannot be queried.
- This scan is on the master node of the instance and will affect instance performance.
- Perform expired key scans during off-peak hours to avoid 100% CPU usage.

#### Scanning and Deleting Expired Keys in a DCS Redis Instance

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

- **Step 2** Click in the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of a DCS Redis instance.
- Step 5 Choose Analysis and Diagnosis > Cache Analysis.
- **Step 6** On the **Expired Key Scan** tab page, scan for expired keys and release them.

The keyspace will be scanned to release the memory used by expired keys that were not released due to the lazy free mechanism.

- Click **Start Analysis** to manually scan expired keys with preset parameters (number of keys to iterate: 100; scan timeout: 360 minutes).
- Enable Scheduled to schedule automatic scans at a specified time. For details about how to configure automatic scans, see Table 10-3 and Automated Scan Performance and Suggestions.

| Parameter | Description   |
|-----------|---|
| Start At  | The first scan can only start after the current time.   |
|           | Format: MM/DD/YYYY hh:mm:ss   |
| Interval  | Interval between scans.   |
|           | <ul> <li>If the previous scan is not complete when the start time<br/>arrives, the upcoming scan will be skipped.</li> </ul>  |
|           | <ul> <li>If the previous scan is complete within five minutes after<br/>the start time, the upcoming scan will not be skipped.</li> </ul>   |
|           | <b>NOTE</b><br>Continuous scans may cause high CPU usage. Set this parameter<br>based on the total number of keys in the instance and the increase<br>of keys. For details, see <b>Automated Scan Performance and</b><br><b>Suggestions</b> . |
|           | Value range: 0–43,200   |
|           | Default value: 1440   |
|           | Unit: minute  |

 Table 10-3 Parameters for scheduling automatic scans

| Parameter       | Description   |  |
|-----------------|---|--|
| Timeout         | This parameter is used to prevent scanning timeout due to<br>unknown reasons. If scanning times out due to unknown<br>reasons, subsequent scheduled tasks cannot be executed.<br>After the specified timeout elapses, a failure message is<br>returned and the next scan will be performed.   |  |
|                 | • Set the timeout to at least twice the interval.   |  |
|                 | • You can set a value based on the time taken in previous scans and the maximum timeout that can be tolerated in the application scenario.  |  |
|                 | Value range: 1–86,400   |  |
|                 | Default value: 2880   |  |
|                 | Unit: minute  |  |
| Keys to Iterate | The <b>SCAN</b> command is used to iterate the keys in the current database. The <b>COUNT</b> option is used to let the user tell the iteration command how many elements should be returned from the dataset in each iteration. For details, see the <b>description of the SCAN command</b> . Iterative scanning can reduce the risks of slowing down Redis when a large number of keys are scanned at a time. |  |
|                 | For example, if there are 10 million keys in Redis and the number of keys to iterate is set to 1000, a full scan will be complete after 10,000 iterations.  |  |
|                 | Value range: 10–1,000   |  |
|                 | Default value: 50   |  |
|                 | Unit: number  |  |

**Step 7** After an expired key scan task is submitted, a task record is generated for each expired key scan. You can view the task ID, status, scan mode, start time, and end time.

Figure 10-4 Expired key scan tasks

| Big Key Analysis Hot Key Analysis Expired Key Sc   | an             |        |                                 |                                 |   |
|--|----------------|--------|---------------------------------|---------------------------------|---|
| () The keyspace will be scanned to release the memory used by expired keys that were not released due to the lazy free mechanism. This scan is on the master node of the instance and will affect instance performance. Learn more |                |        |                                 |                                 |   |
| Scheduled Disabled   |                |        |                                 |                                 |   |
| Start Analysis   |                |        |                                 |                                 |   |
| Q Select a property or enter a keyword.  |                |        |                                 |                                 | 0 |
| Task ID 🕀  | Status \ominus | Mode 😔 | Started 🕀                       | Completed 🕀                     |   |
| 54a4bc76-21ec-478f-a295-8d258c55ec7f   | Successful     | Manual | Feb 25, 2025 11:15:04 GMT+08:00 | Feb 25, 2025 11:15:07 GMT+08:00 |   |

#### **NOTE**

The scan fails in the following scenarios:

- An exception occurred.
- There are too many keys, resulting in a timeout. Some keys have already been deleted before the timeout.

----End

#### Automated Scan Performance and Suggestions

#### Performance

- The SCAN command is executed at the data plane every 5 ms, that is, 200 times per second. If Keys to Iterate is set to 10, 50, 100, or 1000, 2000, 10,000, 20,000, or 200,000 keys are scanned per second.
- The larger the number of keys scanned per second, the higher the CPU usage.

#### **Reference test**

A master/standby instance is scanned. There are 10 million keys that will not expire and 5 million keys that will expire. The expiration time is 1 to 10 seconds. A full scan is executed.

#### **NOTE**

The following test results are for reference only. They may vary depending on the site environment and network fluctuation.

- Natural deletion: 10,000 expired keys are deleted per second. It takes 8 minutes to delete 5 million expired keys. The CPU usage is about 5%.
- **Keys to Iterate** set to **10**: The scanning takes 125 minutes (15 million/ 2000/60 seconds) and the CPU usage is about 8%.
- **Keys to Iterate** set to **50**: The scanning takes 25 minutes (15 million/ 10,000/60 seconds) and the CPU usage is about 10%.
- **Keys to Iterate** set to **100**: The scanning takes 12.5 minutes (15 million/ 20,000/60 seconds) and the CPU usage is about 20%.
- **Keys to Iterate** set to **1000**: The scanning takes 1.25 minutes (15 million/ 200,000/60 seconds) and the CPU usage is about 25%.

#### **Configuration suggestions**

- You can configure the number of keys to be scanned and the scanning interval based on the total number of keys and the increase in the number of keys in the instance.
- In the reference test with 15 million keys and **Keys to Iterate** set to **10**, the scanning takes about 125 minutes. In this case, set the scan interval to more than 4 hours.
- If you want to accelerate the scanning, set **Keys to Iterate** to **100**. It takes about 12.5 minutes to complete the scanning. Therefore, set the scan interval to more than 30 minutes.
- The larger the number of keys to iterate, the faster the scanning, and the higher the CPU usage. There is a trade-off between time and CPU usage.
- If the number of expired keys does not increase rapidly, you can scan expired keys once a day.

#### **NOTE**

Start scanning during off-peak hours. Set the interval to one day and the timeout to two days.

#### **10.3 Analyzing Redis Backup Offline**

The offline key analysis function on the DCS console analyzes the backup of a specific instance node. The analysis covers the top 100 big keys, top 50 keys with the most prefixes of each data type, and the memory usage and number of keys of each data type.

#### **Notes and Constraints**

- Currently, this function is disabled by default. To use it, contact customer service.
- This function is available for DCS Redis 4.0, 5.0, and 6.0 instances.
- Only RDB backups of a single node can by analyzed at a time.
- Existing backups cannot be analyzed after the following changes:
  - Instance scale-in/down
  - Cluster instance scale-out
  - Instance type change (excluding changing from master/standby to read/ write splitting)

#### Procedure

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

- **Step 2** Click <sup>(2)</sup> in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of a DCS instance.
- **Step 5** Choose Analysis and Diagnosis > Cache Analysis > Offline Key Analysis.
- **Step 6** Click **Start Analysis** and select an instance node.

#### Figure 10-5 Selecting a node

| Big Key Analysis Hot Key Analysis Expired Key Scan | Offline Key Analysis  |
|--|---|
| Start Analysis Delete                              | Start Analysis ×  |
| Q. Select a property or enter a keyword.           |   |
| Task ID 😔  | Node  |
|  | (master)192.168.145.40 ~ G  |
|  | Analyzed Backup File  |
|  | Now Archived  |
|  | 1. The new backup is temporary and cannot restore instance data     2. Backup creation deteriorates instance performance. Analyze the node during off-peak hours. |
|  | Cancal OK   |

**Step 7** Specify how to analyze it: Create a backup to analyze or select an archived one.

- New: Create a backup file for the selected node for analysis. (The new backup is used only once in this analysis, and is not recorded in **Backing Up or Restoring Instance Data**).
- Archived: Select a backup file in RDB format from the historical backup records.

#### D NOTE

Analyzing a master node using a new backup may deteriorate instance performance. You are advised to perform it during off-peak hours or analyze a standby node.

**Step 8** Click **OK**. When the analysis task is in the **Successful** state, the analysis is complete.

To download or delete an analysis task, click **Download** or **Delete** on the right of the analysis task. To batch delete analysis tasks, select the tasks and click **Delete** above the list.

Step 9 Click the task ID to view the key analysis result.

A key analysis result covers the basic information, node information, top 100 big keys, top 50 keys with the most prefixes of each data type, and memory usage and quantity distribution of keys of each data type.

#### Figure 10-6 Key analysis result



----End
# **10.4 Diagnosing a DCS Redis Instance**

If a fault or performance issue occurs, you can ask DCS to diagnose your instance to learn about the cause and impact of the issue and how to handle it.

## Notes and Constraints

- DCS Redis 3.0 and Memcached instances do not support diagnosis.
- New instances can be diagnosed 10 minutes after they are successfully created.
- Instance diagnosis may fail during specification modification.

## **Diagnosing a DCS Redis Instance**

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of a DCS Redis instance.
- **Step 5** Choose Analysis and Diagnosis > Instance Diagnosis.
- **Step 6** Specify the tested object and time range, and click **Start Diagnosis**.
  - **Tested Object**: You can select a single node or all nodes.
  - **Range**: You can specify up to 10 minutes before a point in time in the last 7 days.

The data within 10 minutes before the specified time will be diagnosed as shown below.

#### Figure 10-7 Specifying the tested object and time range

| Tested Objects   | -Select nodes to test-   |                                |           | ~                      |   |          |                    |
|------------------|--------------------------|--------------------------------|-----------|------------------------|---|----------|--------------------|
| Range            | - 10 + mi                | n before Dec 09, 2024 16:03:06 |           |                        |   |          |                    |
|                  |                          |                                |           |                        |   |          |                    |
| Test History     |                          |                                |           |                        |   |          |                    |
| Start Diagnosis  | Delete                   |                                |           |                        |   |          |                    |
| Q. Select a prop | erty or enter a keyword. |                                |           |                        |   |          | 0                  |
| Tested 🖯         |                          | Test ID \ominus                | Status 😔  | Tested Objects \ominus | Range 🕀   | Result 🖯 | Operation          |
| Dec 09, 20       | 24 16:03:15 GMT+08:00    | a8b0ef56-ba68-4c99-aa8d-b234   | Completed | All nodes              | Dec 09, 2024 15:53:00 GMT+08:00 - Dec 09, 2024 16:03:00 GMT+08:00 | Normal   | View Report Delete |
| Total Records: 1 | 10 🗸 < 1                 | >                              |           |                        |   |          |                    |

**Step 7** After the diagnosis is complete, you can view the result in the **Test History** list. If the result is abnormal, click **View Report** for details.

In the report, you can view the cause and impact of abnormal items and suggestions for handling them.

----End

# **10.5 Viewing Slow Queries of a DCS Redis Instance**

Redis logs queries that exceed a specified execution time. You can view the slow logs on the DCS console to identify performance issues.

For details about the commands, visit the **Redis official website**.

Configure slow queries with the following parameters:

- **slowlog-log-slower-than**: The maximum time allowed, in microseconds, for command execution. If this threshold is exceeded, Redis will log the command. The default value is **10,000**. That is, if command execution exceeds 10 ms, the command will be logged.
- **slowlog-max-len**: The number of slow queries in a record. The default value is **128**, which means a maximum of 128 latest slow queries can be displayed.

The following parameters are available only in the CN East-Shanghai2 and CN South-Guangzhou regions.

- **proxy-slowlog-log-slower-than**: The maximum time allowed, in microseconds, for command execution. If this threshold is exceeded, Redis will log the command. The default value is **256,000**. That is, if command execution exceeds 256 ms, the command will be logged.
- **proxy-slowlog-max-len**: The number of slow queries in a record. The default value is **128**, which means a maximum of 128 latest slow queries can be displayed.

For details about the configuration parameters, see **Modifying Configuration Parameters of a DCS Instance**.

## **Notes and Constraints**

- You can view the slow queries of a Proxy Cluster DCS Redis 3.0 instance only if the instance is created after October 14, 2019.
- Currently, you can view slow queries in the last seven days.
- After restarting an instance, slow queries before the restart cannot be viewed.
- Currently, slow queries of only Proxy Cluster and read/write splitting instances in the CN North-Beijing4, CN East-Shanghai1, CN East-Shanghai2, and CN South-Guangzhou regions contain **Proxy** and **Redis Server** categories for proxy and Redis instance node records.
  - For Proxy Cluster and read/write splitting instances created before August 2024, and the proxies are not upgraded, contact customer service to upgrade proxies. Otherwise, the slow queries under **Proxy** are always blank.
  - The real client IP address is in the Client IP Address column of the slow query list of an instance with client IP pass-through enabled. For Proxy Cluster and read/write splitting instances, the real client IP address is in the Proxy column.

## Procedure

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

- **Step 2** Click Sin the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the name of a DCS instance.
- **Step 5** Choose **Analysis and Diagnosis** > **Slow Queries**.
- **Step 6** Select a start date and an end date and click the refresh icon to view slow queries within the specified period. For details about the commands, visit the **Redis official website**.

To filter slow queries, click the filtering bar, select a property or enter a keyword.

Figure 10-8 Slow queries of an instance

| ( Lepon - )                     |   |                     |          |                 |                   |                          |      |
|---------------------------------|---|---------------------|----------|-----------------|-------------------|--------------------------|------|
| Last week 🗸 🗸                   | Q. Select a property or enter a keyword |                     |          |                 |                   |                          | Q () |
| Executed 🕀                      | Duration (ms)                           | Shard Name 😔 Role ( | Username | ⊖ Database ID € | Client IP Address | Slow Query 🕀             |      |
| Jun 20, 2024 18:50:51 GMT+08:00 | 13.229                                  | group-1 Master      | default  | 0               | 192. 6            | INFO all                 |      |
| Jun 20, 2024 02:29:24 GMT+08:00 | 11.891                                  | group-1 Master      | default  | 0               | 192. 2            | 4 REPLCONF ACK 346829627 |      |
| Jun 19, 2024 16:47:09 GMT+08:00 | 27.069                                  | group-1 Master      | default  | 0               | 192. 3            | cluster nodes            |      |
| Jun 19, 2024 11:28:03 GMT+08:00 | 15.539                                  | group-1 Master      | default  | 0               | 192. 1            | 01 SLOWLOG GET 10        |      |

Step 7 To download slow queries, choose Export > Export all data to an XLSX file or Export selected data to an XLSX file.

----End

# **10.6 Viewing Redis Run Logs**

Run logs of a Redis instance can be queried on the DCS console. Logs of a specified time can be collected into the **redis.log** file, and downloaded to the local.

Instance running exceptions include AOF rewrites, configuration modifications, critical operations, and master/standby switchovers.

## Notes and Constraints

- This function is supported by DCS Redis 4.0 instances and later.
- The logs are retained for seven days, and are automatically deleted later.
- A maximum of seven days of run logs can be queried for a Redis instance.

#### Procedure

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

**Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.

**Step 3** In the navigation pane, choose **Cache Manager**.

- **Step 4** Click a DCS instance.
- Step 5 Click the Run Logs tab.
- **Step 6** Click **Collect Logs**, specify the collection period, and click **OK**.

If the instance is the master/standby, read/write splitting, or cluster type, you can specify the shard and replica whose run logs you want to collect. If the instance is the single-node type, logs of the only node of the instance will be collected.

A log file contains logs of one day. For example, if you select last 3 days, three log files will be generated.

**Step 7** After the log file is successfully collected, click **Download** to download it.

**NOTE** 

The Redis kernel generates few logs, so your selected period may contain no logs.

----End

# **10.7 Viewing Audit Logs of a DCS Redis Instance**

Command audit logs on the DCS console record client operations on DCS. The storage, query, and analysis of audit logs are provided by Log Tank Service (LTS).

Currently, only **Proxy Cluster DCS Redis 4.0 and later** instances in the CN Beijing4 & Shanghai1–2 & Guangzhou regions support command audit logs.

If audit logs are still not displayed on the console, contact customer service to upgrade your proxies.

Enabling audit logging will create a log group, log stream, and dashboard in LTS. Fees are generated based on the log volume. For details, see LTS pricing details.

#### NOTICE

• Enabling audit logging will restart all proxy nodes. Ensure that the client can re-connect.

Without capacity expansion or node migration since the last enabling, reenabling audit logging will not restart proxy nodes.

• Enabling audit logging may deteriorate DCS instance performance or cause some logs to be lost if the write traffic and QPS are too heavy.

## **Notes and Constraints**

- Ensure that you have permissions to create log groups and log streams in LTS.
- Enable audit logging for new instances 10 minutes after they are successfully created.
- Audit logging will be automatically disabled when you scale the instance or migrate nodes. To use this function, you need to enable it again.

• By default, audit logs only record write operations.

To record read operations, add custom commands to parameter **audit-logcustomer-command-list** by referring to **Modifying Configuration Parameters of a DCS Instance**.

• After enabling audit logging, you can change the log retention period (one day by default) on the LTS console. For details, see **Changing the log retention period**.

# Viewing Audit Logs of a DCS Redis Instance

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

- **Step 2** Click Sin the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click a DCS instance to go to the details page.
- Step 5 Choose Logs > Audit Logs.
- Step 6 Click Enable Audit Logging if required.
- **Step 7** The corresponding log group and log stream are created in LTS. Write commands are reported to LTS through the proxy nodes.
- **Step 8** View the audit logs, as shown in the following figure.



#### Figure 10-9 Audit logs

----End

# **More Operations**

• Disabling audit logging

To disable the function, click **Disable Audit Logging** in the upper right corner. After the function is disabled, commands will not be recorded.

#### **NOTE**

- If some logs are not yet reported by the time you disable audit logging, they continue to be reported.
- After audit logging is disabled, the log group and log stream on LTS will be retained and will not generate fees.
- You can manually delete log groups and log streams on the LTS console.
- Changing the log retention period

In the log group list of the LTS console, click **Modify** in the **Operation** column of the desired log group to change the log retention period.

Figure 10-10 Modifying the log retention period or deleting a log group

| Log Groups   |        |                  |                  |                           |
|--|--------|------------------|------------------|---------------------------|
| Create Log Group Q Click here to choose a filter condition |        |                  |                  | Ó Ŧ                       |
| Log Group Name   | Remark | Enterprise Proje | Log Streams Tags | Operation                 |
| DCS-log-group-1731897660774                                |        | default          | 1                | Modify Delete More $\vee$ |

# **11** Migrating Instance Data

# **11.1 DCS Data Migration Overview**

The DCS console supports online and backup (file) migration with intuitive operations. Incremental data can be migrated online.

• Online migration is suitable when the source Redis instance supports the **SYNC** and **PSYNC** commands. Data in the source Redis instance can be migrated in full or incrementally to the target instance.

During online migration, the **PSYNC** command is delivered to the source address. For details about how this works, see the **replication explanation**. This command will cause a fork operation at the source end, which affects latency. For details about the impact scope, see the **Redis official website**.

 Backup migration is suitable when the source and target Redis instances are not connected, and the source Redis instance does not support the SYNC and PSYNC commands. To migrate data, import your backup files to OBS, and DCS will read data from OBS and migrate the data to the target DCS Redis instance. Alternatively, you can import the backup files directly to the DCS instance.

Users can customize migration solutions as required based on specific Redis environment or scenarios. The data volume, source Redis deployment, and network bandwidth affect migration duration. The actual duration depends.

Before migrating an instance, analyze the cache commands (reference: **Command Compatibility**) used by your service systems and verify the commands one by one during the drill phase. Contact customer service as needed.

#### NOTICE

- Currently, the data migration function is free of charge in the OBT. You will be notified when data migration starts to be charged.
- As an important and stringent task, data migration requires high accuracy and timeliness, which depends on specific services and operations.
- Cases provided in this document are for reference only. Consider your needs during actual migration.
- Some commands in this document contain the instance password, which means the passwords are recorded. Ensure that the passwords are not disclosed and clear operation records in a timely manner.
- DCS for Redis 3.0 is no longer provided. You can use DCS for Redis 4.0 or later.

# **DCS Data Migration Modes**

#### **NOTE**

- DCS for Redis refers to Redis instances provided by Huawei Cloud DCS.
- **Self-hosted Redis** refers to self-hosted Redis on the cloud, from other cloud vendors, or in on-premises data centers.
- √: Supported. ×: Not supported.
- You can migrate data online in full or incrementally from **other cloud Redis** to **DCS for Redis** if they are connected and the **SYNC** and **PSYNC** commands can be run on the source Redis. However, some instances provided by other cloud vendors may fail to be migrated online. In this case, migrate data through backup import or use other migration schemes. For more information, see **Migration Solution Notes**.

| Migration                   | Source  | Target: DCS   |               |               |  |  |
|-----------------------------|---|---|---------------|---------------|--|--|
| Mode                        |   | Single-node, read/<br>write splitting, or<br>master/standby | Proxy Cluster | Redis Cluster |  |  |
| Importing                   | AOF file  | $\checkmark$  | $\checkmark$  | $\checkmark$  |  |  |
| backup<br>files             | RDB file  | $\checkmark$  | $\checkmark$  | $\checkmark$  |  |  |
| Migrating<br>data<br>online | DCS for Redis:<br>Single-node,<br>read/write<br>splitting, or<br>master/standby | $\checkmark$  | $\checkmark$  | $\checkmark$  |  |  |
|                             | DCS for Redis:<br>Proxy Cluster   | $\checkmark$  | $\checkmark$  | $\checkmark$  |  |  |
|                             | DCS for Redis:<br>Redis Cluster   | $\checkmark$  | $\checkmark$  | $\checkmark$  |  |  |
|                             | Self-hosted Redis   | $\checkmark$  | $\checkmark$  | $\checkmark$  |  |  |

Table 11-1 DCS data migration modes

| Other cloud | $\checkmark$ | $\checkmark$ | $\checkmark$ |
|-------------|--------------|--------------|--------------|
| Redis       |              |              |              |

# **Migration Process**



1. Evaluation

Collect the following information about the cached data to be migrated (based on Table 11-2):

- Number of instances
- Number of databases (DBs) configured for each instance
- Number of keys in each DB
- DBs used for your services
- Space occupied by each instance
- Redis version
- Redis instance type
- Mapping relationships between your services and instances

#### **NOTE**

- The **info keyspace** command can be used on accessed source Redis to check whether databases have available data and the key quantity in databases. The key quantity in each database can be used for migration verification.
- The **info memory** command can be used on accessed source Redis to check the source Redis data volume by the returned **used\_memory\_human** value. Check whether the following resources are sufficient: available ECS space, target instance flavor, and remaining memory (≥ source Redis data).

Plan the following information about target DCS instances based on the collected information:

- Number of DCS instances
- Memory size of the DCS instance (≥ that of source Redis)
- Version of the DCS instance ( $\geq$  that of source Redis)
- DCS instance type

- Virtual Private Clouds (VPCs), security groups, and subnets, and security groups, to which the instances and services belong

#### 2. Preparation

After completing the evaluation, prepare the following items:

a. Mobile storage devices

These devices are used to copy and transfer data in case of network disconnection (in scenarios with data centers of enterprises).

b. Network resources

Create VPCs and subnets based on service planning.

c. Server resources

**Purchase ECSs** to bear Redis clients. The ECSs are used to export or import cached data.

Recommended ECS specifications are 8 vCPUs | 16 GB or higher.

d. DCS instances

**Purchase DCS instances** based on the migration planning. If the number of instances exceeds the default quota, submit a service ticket or contact customer service.

e. Related tools

Install the FTP tool and SSH tool.

f. Information to be collected

Collect the contact information of people involved in the migration, ECS login credentials, cache instance information, and DB information.

g. Overall migration plan

Formulate the overall migration plan, including the personnel arrangement, drill, migration, verification, service switchover, and rollback solutions.

Break down each solution into executable operations and set milestones to mark the end of tasks.

3. Drill

The drill phase aims to:

- a. Verify the feasibility of the migration tools and migration process.
- b. Discover problems that may occur during migration and make effective improvements.
- c. Evaluate the time required for migration.
- d. Optimize the migration steps and verify the feasibility of concurrent implementation of some tasks to improve migration efficiency.
- 4. Backup

Before migration, back up related data, including but not limited to cached data and Redis configuration files, in case of emergency.

5. Migration

After conducting one or two rounds of migration drill and solving problems found in the drill, start data migration.

Break down the migration process into executable steps with specific start and end confirmation actions.

6. Data verification

Check the following items:

- The key distribution of each DB is consistent with the original or expected distribution.
- Main keys.
- Expiration time of keys.
- Whether instances can be normally backed up and restored.
- 7. Service switchover
  - a. After the data migration and verification, use the new instances for your services.
  - b. If DB IDs are changed, modify the ID configurations for your services.
  - c. If your services are migrated from data centers or cloud platforms provided by other vendors to Huawei Cloud as a whole, services and cached data can be migrated concurrently.
- 8. Service verification
  - a. Verify the connectivity between your service applications and DCS instances.
  - b. Verify whether cached data can be normally added, deleted, modified, and queried.
  - c. If possible, perform pressure tests to ensure that the performance satisfies the peak service pressure.
- 9. Rollback

If your services are unavailable after the data migration because unexpected problems occur and cannot be solved in the short term, roll back your services.

Since source Redis data still exists, you only need to roll back your services and use the source Redis instances again.

After the rollback, you can continue to restart from the drill or even preparation phase to solve the problems.

# Information to be collected for the migration

The following table lists the information to be collected in the evaluation and preparation phases.

| Migration<br>Source   | ltem   | Description  |
|---|--|--|
| Source<br>Redis   | Source Redis IP<br>address                         | -  |
| (List the<br>informatio<br>n about all<br>instances   | Redis instance<br>password (if<br>any)             | -  |
| to be<br>migrated.)   | Total data<br>volume                               | Run the <b>info memory</b> command and refer to the <b>used_memory_human</b> value to obtain the total data volume.  |
|   |  | Used to evaluate whether the migration<br>solution, DCS instance specifications, and<br>available disk space of ECSs meet requirements,<br>and to estimate the time required for migration<br>(service interruption duration). |
|   | IDs of DBs with data                               | Obtained by running the <b>info keyspace</b> command.  |
|   |  | Used to check whether the migration involves<br>multiple DBs and non-AOF files. Some open-<br>source tools can export and import data of only<br>one DB at a time.   |
|   |  | For DCS instances, the single-node and master/<br>standby types provide 256 DBs (DB 0 to DB<br>255), and the cluster type provides only one DB<br>by default.  |
|   | Number of keys<br>in each DB                       | Used to verify the data integrity after migration.   |
|   | Data type  | The Cloud Data Migration (CDM) service<br>supports two data formats: hash and string. If<br>the source data contains data in other formats<br>such as list and set, use a third-party migration<br>tool.                       |
| Huawei<br>Cloud ECS<br>If a large   | EIP  | Select ECSs that can communicate with DCS instances for data import to ensure network stability.   |
| number of<br>instances<br>are to be<br>migrated,<br>prepare<br>multiple<br>ECSs for<br>concurrent<br>migration. |  | Configure high-specification bandwidth to improve data transmission efficiency.  |
|   | Login<br>credentials<br>(username and<br>password) | -  |

| Гable | 11-2 | Information | to be | collected | for the | e migration |
|-------|------|-------------|-------|-----------|---------|-------------|
|-------|------|-------------|-------|-----------|---------|-------------|

| Migration<br>Source   | Item  | Description  |
|---|---|--|
|   | CPU and<br>memory                                     | Some migration tools support concurrent<br>import through multiple threads. High-<br>specification ECSs help improve import<br>efficiency.   |
|   | Available disk<br>space                               | Sufficient available disk space needs to be reserved on the ECSs to store compressed files and decompressed cached data files.   |
|   |   | Note: To improve data transmission efficiency,<br>compress large-size data files before<br>transmitting them to ECSs.  |
| DCS<br>instances<br>(Select                                       | Instance<br>connection<br>address                     | -  |
| appropriate<br>instance<br>specificatio                           | Instance<br>connection port                           | -  |
| ns and<br>quantities  | Instance<br>password                                  | -  |
| the   | Instance type   | -  |
| number of<br>source<br>Redis<br>instances<br>and data<br>volume.) | Instance<br>specifications<br>and available<br>memory | -  |
| Network<br>configurati<br>ons                                     | VPC   | Plan VPCs in advance to ensure that your service applications and DCS instances are in same VPCs.  |
|   | Subnet  | -  |
|   | Whitelist or<br>security group                        | DCS Redis 3.0, and 4.0 and later enterprise<br>edition instances are deployed in different<br>modes. Therefore, the access control methods<br>vary. You can control access to your DCS<br>instances by setting security groups or<br>whitelists. For details, see How Do I Configure<br>a Security Group? or Managing IP Address<br>Whitelist. |
| -   | -   | Other configurations.  |

# **11.2 Migration Solution Notes**

# **Migration Tools**

| Tool/<br>Command/<br>Service | Feature   | Description   |
|------------------------------|---|---|
| DCS console                  | Supports online<br>migration (in full or<br>incrementally) and<br>backup migration (by<br>importing backup files)<br>with intuitive operations.   | <ul> <li>Backup migration is suitable<br/>when the source and target<br/>Redis instances are not<br/>connected, and the source Redis<br/>instance does not support the<br/>SYNC and PSYNC commands. To<br/>migrate data, import your<br/>backup files to OBS, and DCS<br/>will read data from OBS and<br/>migrate the data to the target<br/>DCS Redis instance.</li> <li>Online migration is suitable</li> </ul> |
|                              |   | when the source Redis instance<br>supports the <b>SYNC</b> and <b>PSYNC</b><br>commands. Data in the source<br>Redis instance can be migrated<br>in full or incrementally to the<br>target instance.  |
| redis-cli                    | <ul> <li>The Redis command<br/>line interface (CLI),<br/>which can be used to<br/>export data as an<br/>RDB file or import the<br/>AOF file (that is, all<br/>DBs) of an instance.</li> <li>An AOF file is large<br/>file containing a full<br/>set of data change<br/>commands.</li> </ul> | -   |
| Rump                         | Supports online<br>migration between DBs<br>of an instance or<br>between DBs of different<br>instances.   | Rump does not support incremental<br>migration.<br>Stop services before migrating data.<br>Otherwise, keys might be lost. For<br>details, see <b>Online Migration from</b><br><b>Another Cloud Using Rump</b> .   |

| <b>Table</b> | 11-3 | Comparing | Redis | migration | tools |
|--------------|------|-----------|-------|-----------|-------|
|              |      |           |       |           |       |

| Tool/<br>Command/<br>Service              | Feature  | Description  |
|---|--|--|
| Redis-shake                               | An open-source tool that supports both online and offline migration. | redis-shake is suitable for migrating<br>Redis Cluster data. |
| Self-<br>developed<br>migration<br>script | Flexible and can be adjusted as required.                            | -  |

# **Migration Schemes**

| Table 11-4 M | igration Schemes |
|--------------|------------------|
|--------------|------------------|

| Scenario   | Tool        | Use Case   | Description   |
|--|-------------|--|---|
| Migration<br>between<br>Huawei<br>Cloud DCS<br>instances | DCS console | <ul> <li>To migrate instances in a region and of an account, see Online Migration Between Instances.</li> <li>To migrate instances in different regions or accounts, see Backup Import Between DCS Redis Instances.</li> </ul> | Attempts to<br>migrate data<br>from a later-<br>version Redis<br>instance to an<br>earlier-version<br>Redis instance<br>are not<br>recommended<br>because they<br>will fail due to<br>data<br>compatibility<br>issues between<br>different Redis<br>versions. |

| Scenario  | Tool        | Use Case  | Description  |
|---|-------------|---|--|
| From self-<br>hosted<br>Redis to<br>DCS<br>NOTE<br>Self-<br>hosted<br>Redis<br>refers to<br>self-<br>hosted<br>Redis on<br>Huawei<br>Cloud, in<br>another | DCS console | <ul> <li>If the network between your self-hosted Redis instance and the DCS Redis instance is connected, follow to the instructions in Online Migration Between Instances.</li> <li>If the network between your self-hosted Redis instance and the DCS Redis instance is not connected, follow to the instructions in Self-Hosted Redis Migration with Backup Files.</li> </ul> | -  |
| in on-<br>premises<br>data<br>centers.  | redis-cli   | <ul> <li>Self-Hosted Redis<br/>Migration with redis-cli<br/>(AOF)</li> <li>Self-Hosted Redis<br/>Migration with redis-cli<br/>(RDB)</li> </ul>  | -  |
|   | redis-shake | <ul> <li>Self-Hosted Redis Cluster<br/>Migration with redis-shake<br/>(Online)</li> <li>Self-Hosted Redis Cluster<br/>Migration with redis-shake<br/>(RDB)</li> </ul>   | -  |
| From<br>another<br>cloud to<br>DCS  | DCS console | <ul> <li>If the SYNC and PSYNC commands are not disabled for the Redis service provided by another cloud, follow the instructions in Migrating Redis from Another Cloud Online.</li> <li>If the SYNC and PSYNC commands are disabled for the Redis service provided by another cloud, follow the instructions in Backup Import from Another Cloud.</li> </ul>                   | If online<br>migration is<br>required, contact<br>the O&M<br>personnel of<br>another cloud to<br>enable the <b>SYNC</b><br>and <b>PSYNC</b><br>commands. |
|   | Rump        | Online Migration from<br>Another Cloud Using Rump   | -  |

| Scenario | Tool        | Use Case  | Description |
|----------|-------------|---|-------------|
|          | Redis-shake | Backup Import from Another<br>Cloud Using redis-shake<br>Migrating from Another<br>Cloud Online Using redis-<br>shake | -           |

# **11.3 Migrating Data Between DCS Instances**

# **11.3.1 Online Migration Between Instances**

If the source and target instances are interconnected and the **SYNC** and **PSYNC** commands are supported by the source instance, data can be migrated online in full or incrementally from the source to the target.

#### 

During online migration, data is essentially synchronized in full to a new replica. Therefore, perform online migration during low-demand hours. Otherwise, source instance CPU usage may surge and latency may increase.

## Notes and Constraints

- You cannot use public networks for online migration.
- Before migrating data, read through Migration Solution Notes to learn about the DCS data migration function and select an appropriate target instance.
- Migrating a later Redis instance to an earlier one may fail.
- For earlier instances whose passwords contain single quotation marks ('), modify the password for online migration or try other methods.
- By default, a Proxy Cluster instance has only one database (DB0). Before you
  migrate data from a multi-DB single-node or master/standby instance to a
  Proxy Cluster instance, check whether any data exists on databases other than
  DB0. If yes, enable multi-DB for the Proxy Cluster instance by referring to
  Enabling Multi-DB.
- By default, a Redis Cluster instance has only one DB (DB0). Before you migrate data from a multi-DB single-node or master/standby instance to a Redis Cluster instance, check whether any data exists on databases other than DB0. To ensure that the migration succeeds, move all data to DB0 by referring to Online Migration from Another Cloud Using Rump.
- During online migration, you are advised to set **repl-timeout** on the source instance to 300s and **client-output-buffer-slave-hard-limit** and **client-output-buffer-slave-soft-limit** to 20% of the maximum memory of the instance.
- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

# Prerequisites

- If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data.

If the data exists on the target instance, the replicated data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.

# **Creating an Online Migration Task**

#### NOTICE

Only when the online migration task and the source Redis are under an account and in a region, the **SYNC** and **PSYNC** commands of the source Redis are allowed. Therefore, create one task under the same account in the same region as the source.

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

If the source and target Redis are under different accounts, use the source account to log in to DCS.

- **Step 2** Click Sin the upper left corner of the console and select the region where your **source** instance is located.
- **Step 3** In the navigation pane, choose **Data Migration**. The migration task list is displayed.
- Step 4 Click Create Online Migration Task.
- **Step 5** Enter the task name and description.

The task name must start with a letter, contain 4 to 64 characters, and contain only letters, digits, hyphens (-), and underscores (\_).

**Step 6** Configure the VPC, subnet, and security group for the migration task.

**NOTE** 

- Use the VPC of the source or target Redis.
- The online migration task uses a tenant IP address (**Migration ECS** displayed on the **Basic Information** page of the task.) If a whitelist is configured for the source or target instance, add the migration IP address to the whitelist or disable the whitelist.
- To allow the VM used by the migration task to access the source and target instances, set an outbound rule for the task's security group to allow traffic through the IP addresses and ports of the source and target instances. By default, all outbound traffic is allowed.

----End

# **Checking the Network**

**Step 1** Check whether the source Redis instance, the target Redis instance, and the migration task are configured with the same VPC.

If yes, go to **Configuring the Online Migration Task**. If no, go to **Step 2**.

**Step 2** Check whether the VPCs configured for the source Redis instance, the target Redis instance, and the migration task are connected to ensure that the VM resource of the migration task can access the source and target Redis instances.

If yes, go to **Configuring the Online Migration Task**. If no, go to **Step 3**.

- **Step 3** Perform the following operations to establish the network.
  - If the source and target Redis instances are in the same DCS region, create a VPC peering connection by referring to VPC Peering Connection.
  - If the source and target Redis instances are in different DCS regions, create a cloud connection by referring to Cloud Connect Getting Started.

----End

## **Configuring the Online Migration Task**

- **Step 1** On the **Online Migration** tab page, click **Configure** in the row containing the online migration task you just created.
- **Step 2** Select a migration type.

Supported migration types are **Full** and **Full + Incremental**, which are described in **Table 11-5**.

To **switch DCS instance IPs** after instance migration, select **Full + Incremental** for the migration type.

| Migration Type     | Description  |
|--------------------|--|
| Full               | Suitable for scenarios where services can be interrupted.<br>Data is migrated at one time. <b>Source instance data</b><br><b>updated during the migration will not be migrated to</b><br><b>the target instance.</b>   |
| Full + incremental | Suitable for scenarios requiring minimal service<br>downtime. The incremental migration parses logs to<br>ensure data consistency between the source and target<br>instances.  |
|                    | Once the migration starts, it remains <b>Migrating</b> until you click <b>Stop</b> in the <b>Operation</b> column. After the migration is stopped, data in the source instance will not be lost, but data will not be written to the target instance. When the transmission network is stable, the delay of incremental migration is within seconds. The actual delay depends on the transmission quality of the network link. |

| Table | 11-5 | Migration | type | description |
|-------|------|-----------|------|-------------|
|-------|------|-----------|------|-------------|

#### **Figure 11-2** Selecting the migration type

| ★ Migration Type | • Full<br>Suitable for scenarios where services can be interrupted. Data is migrated at one time. Source Redis data updated during the migration will not be migrated to the target instance. |
|------------------|---|
|                  | Full + Incremental<br>Suitable for scenarios requiring minimal service downtime. The incremental migration parses logs to ensure data consistency between the source Redis and target Redis.  |

**Step 3** Only if **Migration Type** is set to **Full + Incremental**, you can specify a bandwidth limit.

The data synchronization rate can be kept around the bandwidth limit.

**Step 4** Specify **Auto-Reconnect**. If this option is enabled, automatic reconnections will be performed indefinitely in the case of a network exception.

Full synchronization will be triggered and requires more bandwidth if incremental synchronization becomes unavailable. Exercise caution when enabling this option.

#### Step 5 Configure Source Redis and Target Redis.

1. Set Source Redis Type to Redis in the cloud and add Source Redis Instance.

**NOTE** 

When a DCS instance is used as the source, do not select **Self-hosted Redis**.

- 2. Configure Target Redis Type and Target Redis Instance:
  - If the target Redis and migration task are in a VPC, or across VPCs over a network in a region, set **Target Redis Type** to **Redis in the cloud** and add **Target Redis Instance**.
  - If the target Redis and migration task are in different regions, set Target Redis Type to Self-hosted Redis and add Target Redis Instance. If the target Redis is a Redis Cluster, enter the IP addresses and ports of all masters in the cluster and separate multiple addresses with commas (,). For example: 192.168.1.1:6379,192.168.0.0:6379
- 3. Configure **Source Redis Instance Password** and **Target Redis Instance Password**: If the instance is password-protected, click **Test Connection** to check whether the instance password is correct and whether the network is connected. If the instance is not password-protected, click **Test Connection** directly.

Currently, the users created in **Managing Users** are unavailable here.

4. You can specify the source DB and target DB. For example, if you enter **5** for source DB and **6** for target DB, data in DB5 of the source Redis will be migrated to DB6 of the target Redis. If the source DB is not specified but the target DB is specified, all source data will be migrated to the specified target DB by default. If the target DB is not specified, data will be migrated to the corresponding target DB.

#### **NOTE**

- If the source Redis is multi-DB and the target is single-DB (DB0), either ensure that all source data is in DB0, or specify a source DB and set the target DB to 0.
   Otherwise, migration will fail.
- For details about DB in DCS for Redis, see Does DCS for Redis Support Multi-DB?

#### Step 6 Click Next.

Step 7 Confirm the migration task details and click Submit.

Go back to the data migration task list. After the migration is successful, the task status changes to **Successful**.

If the migration fails, click the migration task and check the log on the **Migration Logs** page.

**NOTE** 

- Once full + incremental migration starts, it remains **Migrating** after full migration.
- To manually stop a migration task, select the check box on the left of the migration task and click **Stop** above the migration task.
- To perform migration again, select the migration tasks which failed or are stopped, and click **Restart** above. If a restarted migration task fails, click **Configure** to configure the task and try again.
- A maximum of 50 online migration tasks can be selected at a time. You can stop, delete, or restart them in batches.

----End

## Verifying the Migration

After the migration is complete, check data integrity in the following way.

- 1. Connect the source Redis and the target Redis. For details, see redis-cli.
- 2. Run the **info keyspace** command on the source and the target Redis to check the values of **keys** and **expires**.

Figure 11-3 Checking instance data

192.100.217:6379> info keyspace
# Keyspace
db0:keys=81869,expires=0,avg\_ttl=0
192.100.00.217:6379>

3. Calculate the differences between the values of **keys** and **expires** of the source Redis and the target Redis. If the differences are the same, the data is complete and the migration is successful.

#### **NOTE**

During full migration, source Redis data updated during the migration will not be migrated to the target instance.

## (Optional) Switching DCS Instance IP Addresses

The prerequisites for switching source and target Redis instance IP addresses are as follows. The target Redis can be accessed automatically on a client after the switch.

#### **Prerequisites:**

• This function is supported by basic edition DCS Redis 4.0 instances and later, **but not by enterprise edition DCS Redis instances**.

- For DCS Redis 3.0 instances, contact customer service to enable the whitelist for Redis 3.0 instance IP switches. The instance IP addresses can be switched only when the source instance is a DCS Redis 3.0 instance and the target instance is a basic edition DCS Redis 4.0 or later instance.
- The IP addresses of a source or target instance with public access enabled cannot be switched.
- Instance IPs can be switched only for the source and target Redis that are single-node, master/standby, read/write splitting, or Proxy Cluster instances.
- Full + Incremental must be selected in Step 2.
- The source and target Redis instance ports must be consistent.

#### NOTICE

- 1. Online migration will stop during the switching.
- 2. Instances will be read-only for one minute and disconnected for several seconds during the switching. When the source is a Redis 3.0 instance, the instance will be read-only for one minute and disconnected for 30 seconds during an IP switch.
- 3. If your application cannot reconnect to Redis or handle exceptions, you may need to restart the application after the IP switching.
- 4. If the source and target instances are in different subnets, the subnet information will be updated after the switching.
- 5. If the source is a master/standby instance, the IP address of the standby node will not be switched. Ensure that this IP address is not used by your applications.
- 6. If your applications use a domain name to connect to Redis, the domain name will be used for the source instance. Select **Yes** for **Switch Domain Name**.
- 7. Ensure that the passwords of the source and target instances are the same. If they are different, verification will fail after the switching.
- 8. If a whitelist is configured for the source instance, ensure that the same whitelist is configured for the target instance before switching IP addresses.
- 9. After the IP addresses of a DCS Redis 3.0 instance are switched, synchronize the security group of the source to the whitelist of the target.
- Step 1 On the Data Migration > Online Migration page, when the migration task status changes to Incremental migration in progress, choose More > Switch IP in the Operation column.
- **Step 2** In the **Switch IP** dialog box, select whether to switch the domain name.

**NOTE** 

- If a Redis domain name is used on the client, switch it or you must modify the domain name on the client.
- If the domain name switch is not selected, only the instance IP addresses will be switched.
- **Step 3** Click **OK**. The IP address switching task is submitted successfully. When the status of the migration task changes to **IP switched**, the IP address switching is complete.

To restore the IPs, choose **More** > **Roll Back IP** in the operation column. The IPs are rolled back when the task is in the **Successful** state.

----End

# 11.3.2 Backup Import Between DCS Redis Instances

You can migrate data between DCS instances by importing backup files.

- If the source Redis and target Redis are in the same region under the same DCS account, and the source Redis is not a single-node instance, see **Importing Backup Data from a Redis Instance**.
- If the source Redis and target Redis are in different regions or under different DCS accounts, or the source Redis is a single-node instance, see Importing Backup Data from an OBS Bucket.

## Notes and Constraints

- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.
- Migration may fail if the target instance uses smaller specifications than its source.

## Prerequisites

- You have successfully backed up the source Redis instance.
  - For **Importing Backup Data from a Redis Instance**, you do not need to download the backup file to the local PC. For details about how to back up data, see **Manually Backing Up a DCS Instance**.
  - For Importing Backup Data from an OBS Bucket, download the backup file to the local PC by referring to Downloading a Backup File.
- You have prepared the target Redis instance. If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.

Redis is backward compatible. The target instance version must be the same as or later than the source instance version.

• Ensure that the target Redis instance has sufficient storage space. You can clear the instance data before the migration. For details, see **Clearing DCS Instance Data**. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.

## Importing Backup Data from a Redis Instance

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

- **Step 2** Click O in the upper left corner of the console and select the region where your source and target instances are located.
- **Step 3** In the navigation pane, choose **Data Migration**. The migration task list is displayed.

#### Step 4 Click Create Backup Import Task.

**Step 5** Enter the task name and description.

The task name must start with a letter, contain 4 to 64 characters, and contain only letters, digits, hyphens (-), and underscores (\_).

Step 6 For source Redis, set Data Source to Redis.

#### Figure 11-4 Selecting a data source (Redis)

| Source Redis            |            |       |
|-------------------------|------------|-------|
| ★ Data Source           | OBS bucket | Redis |
| * Source Redis Instance |            | (+)   |

- **Step 7** For **Source Redis Instance**, select the source instance to be migrated.
- Step 8 You can specify Source DB to migrate data from the specified DB in the source Redis backup file. For example, if you enter 5, only data in DB5 will be migrated. To migrate all databases, do not specify Source DB.
- **Step 9** Enable **Multi-DB Proxy Cluster** if the source Redis is a multi-DB (**multi-db** set to **yes**) Proxy Cluster DCS Redis instance.
- **Step 10** Select the backup task whose data is to be migrated.
- **Step 11** For **Target Redis Instance**, select the DCS Redis instance prepared in **Prerequisites**.
- **Step 12** If the target Redis instance has a password, enter the password and click **Test Connection** to check whether the password is correct. If the instance is not password-protected, click **Test Connection** directly.
- Step 13 For Target DB, you can specify a DB in the target Redis to migrate data to. For example, if you enter 5, data will be migrated to DB5 of the target Redis. If you do not specify a DB, data will be migrated to a DB corresponding to the source DB.

#### **NOTE**

- If the source Redis is multi-DB and the target is single-DB (DB0), either ensure that all source data is in DB0, or specify a source DB and set the target DB to **0**. Otherwise, migration will fail.
- For details about DB in DCS for Redis, see **Does DCS for Redis Support Multi-DB?**

#### Step 14 Click Next.

**Step 15** Confirm the migration task details and click **Submit**.

Go back to the data migration task list. After the migration is successful, the task status changes to **Successful**.

----End

# Importing Backup Data from an OBS Bucket

Simply download the source Redis data and then upload the data to an OBS bucket in the same account and region as the target DCS Redis instance. After you

have created a backup import task, data in the OBS bucket will be read and migrated to the target Redis.

#### **NOTE**

- .aof, .rdb, .zip, and .tar.gz files can be uploaded to OBS buckets. You can directly upload .aof and .rdb files or compress them into .zip or .tar.gz files before uploading.
- To migrate data from a cluster Redis instance, download all backup files and upload all of them to the OBS bucket. Each backup file contains data for a shard of the instance. During the migration, you need to select backup files of all shards.

#### Step 1 Create an OBS bucket in the account and region where the target Redis instance is located. If a qualified OBS bucket is available, you do not need to create one.

When creating an OBS bucket, pay attention to the configuration of the following parameters. For details on how to set other parameters, see **Creating a Bucket**.

Region:

The OBS bucket must be in the same region as the target DCS Redis instance.

• **Default Storage Class**: Select **Standard** or **Infrequent Access**. Do not select **Archive**. Otherwise, the migration will fail.

**Step 2** Upload the backup file to the OBS bucket.

- 1. In the bucket list, click the name of the created bucket.
- 2. In the navigation pane, choose **Objects**.
- 3. On the **Objects** tab page, click **Upload Object**.
- 4. Specify **Storage Class**.

Do not select **Archive**. Otherwise, the migration will fail.

5. Upload the objects.

Drag files or folders to the **Upload Object** area or click **add file**.

A maximum of 100 files can be uploaded at a time. The total size cannot exceed 5 GB.

#### Figure 11-5 Uploading an object

|                        | (Optional) Con                                  | figure Advanced Settings                                    |   |   |      |
|------------------------|---|---|---|---|------|
| Storage Class          | Standard  | Infrequent Access   | Archive   |   |      |
|                        | Optimized for frequent                          | ly accessed (multiple times                                 | per month) data such a                                | as small and essential files that require low later           | icy. |
|                        | If you do not change the bucket creation. Learn | nis setting, your uploaded o<br>n more                      | bjects will be stored usi                             | ng the default storage class you selected during              |      |
| Jpload Object          | The file or fold versions of the                | er you newly upload will ove<br>same file or folder, enable | erwrite any existing file oversioning for the current | or folder with the same name. To keep different<br>nt bucket. |      |
|                        |   | Dras as   | d drap files or folders                               | r add files   |      |
|                        | (A n  | Diag an   | unloaded at a time. Th                                | a total size cannot exceed 5 GB )                             |      |
|                        | (A n  | haximum of 100 files can be                                 | uploaded at a time. Th                                | e total size cannot exceed 5 GB.)                             |      |
| Server-Side Encryption | (A n<br>SSE-KMS                                 | naximum of 100 files can be<br>SSE-OBS                      | uploaded at a time. Th<br>Disable                     | e total size cannot exceed 5 GB.)                             |      |

- 6. Specify Server-Side Encryption. If you enable it, you can select SSE-KMS or SSE-OBS. You can also disable it. For details, see Server-Side Encryption.
- 7. Click Upload.
- **Step 3** Click in the upper left corner and choose **Distributed Cache Service for Redis** under **Middleware** to open the DCS console.
- **Step 4** In the navigation pane, choose **Data Migration**.
- Step 5 Click Create Backup Import Task.
- **Step 6** Enter the task name and description.

The task name must start with a letter, contain 4 to 64 characters, and contain only letters, digits, hyphens (-), and underscores (\_).

**Step 7** In the **Source Redis** area, select **OBS Bucket** for **Data Source** and then select the OBS bucket to which you have uploaded backup files.

Figure 11-6 Selecting a data source (OBS bucket)

| Source Redis  |            |                 |
|---------------|------------|-----------------|
| ★ Data Source | OBS bucket | Redis           |
| ★ OBS Bucket  |            | ✓ Q View Bucket |

- Step 8 You can specify Source DB to migrate data from the specified DB in the source Redis backup file. For example, if you enter 5, only data in DB5 will be migrated. To migrate all databases, do not specify Source DB.
- **Step 9** Enable **Multi-DB Proxy Cluster** if the source Redis is a multi-DB (**multi-db** set to **yes**) Proxy Cluster DCS Redis instance.
- **Step 10** Click **Add Backup** and select the backup files to be migrated.
- Step 11 In the Target Redis area, select the Target Redis Instance prepared in Prerequisites.
- **Step 12** If the target Redis instance has a password, enter the password and click **Test Connection** to check whether the password is correct. If the instance is not password-protected, click **Test Connection** directly.
- Step 13 For Target DB, you can specify a DB in the target Redis to migrate data to. For example, if you enter 5, data will be migrated to DB5 of the target Redis. If you do not specify a DB, data will be migrated to a DB corresponding to the source DB.

**NOTE** 

- If the source Redis is multi-DB and the target is single-DB (DB0), either ensure that all source data is in DB0, or specify a source DB and set the target DB to **0**. Otherwise, migration will fail.
- For details about DB in DCS for Redis, see **Does DCS for Redis Support Multi-DB**?
- Step 14 Click Next.
- **Step 15** Confirm the migration task details and click **Submit**.

Go back to the data migration task list. After the migration is successful, the task status changes to **Successful**.

----End

# 11.4 Migrating Data from Self-Hosted Redis to DCS

# 11.4.1 Migrating Self-Built Redis Online

If the source self-host and target instances are interconnected and the **SYNC** and **PSYNC** commands are supported by the source instance, data can be migrated online in full or incrementally from the source to the target DCS.

#### **NOTE**

During online migration, data is essentially synchronized in full to a new replica. Therefore, perform online migration during low-demand hours. Otherwise, source instance CPU usage may surge and latency may increase.

# **Notes and Constraints**

- If the **SYNC** and **PSYNC** commands are disabled by the source instance, enable them before migrating data. Otherwise, the migration fails.
- You cannot use public networks for online migration.

- During online migration, you are advised to set **repl-timeout** on the source instance to 300s and **client-output-buffer-slave-hard-limit** and **client-output-buffer-slave-soft-limit** to 20% of the maximum memory of the source instance.
- The source must be Redis 3.0 or later.
- For earlier instances whose passwords contain single quotation marks ('), modify the password for online migration or try other methods.
- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

## Prerequisites

- Before migrating data, read through **Migration Solution Notes** to learn about the DCS data migration function and select an appropriate target instance.
- By default, a Proxy Cluster instance has only one database (DB0). Before you
  migrate data from a multi-DB single-node or master/standby instance to a
  Proxy Cluster instance, check whether any data exists on databases other than
  DB0. If yes, enable multi-DB for the Proxy Cluster instance by referring to
  Enabling Multi-DB.
- By default, a Redis Cluster instance has only one DB (DB0). Before you
  migrate data from a multi-DB single-node or master/standby instance to a
  Redis Cluster instance, check whether any data exists on databases other than
  DB0. To ensure that the migration succeeds, move all data to DB0 by referring
  to Online Migration from Another Cloud Using Rump.
- The IP address and port of the source Redis instance has been obtained.
- If a target DCS Redis instance is not available, create one first. For details, see Buying a DCS Redis Instance.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.

## **Creating an Online Migration Task**

- **Step 1** Log in to the DCS console using the account of the target DCS Redis instance.
- **Step 2** Click Sin the upper left corner of the console and select the region where your target instance is located.
- Step 3 In the navigation pane, choose Data Migration.
- Step 4 Click Create Online Migration Task.
- **Step 5** Enter the task name and description.

The task name must start with a letter, contain 4 to 64 characters, and contain only letters, digits, hyphens (-), and underscores (\_).

**Step 6** Configure the VPC, subnet, and security group for the migration task.

#### D NOTE

- Select the same VPC as the target Redis. Ensure that the migration resource can access the target Redis instance.
- The online migration task uses a tenant IP address (**Migration ECS** displayed on the **Basic Information** page of the task.) If a whitelist is configured for the source or target instance, add the migration IP address to the whitelist or disable the whitelist.
- To allow the VM used by the migration task to access the source and target instances, set an outbound rule for the task's security group to allow traffic through the IP addresses and ports of the source and target instances. By default, all outbound traffic is allowed.

----End

## **Checking the Network**

**Step 1** Check whether the source Redis instance, the target Redis instance, and the migration task are configured with the same VPC.

If yes, go to Creating an Online Migration Task. If no, go to Step 2.

**Step 2** Check whether the VPCs configured for the source Redis instance, the target Redis instance, and the migration task are connected to ensure that the VM resource of the migration task can access the source and target Redis instances.

If yes, go to **Configuring the Online Migration Task**. If no, go to **Step 3**.

- **Step 3** Perform the following operations to establish the network.
  - If the VPC of the source and target Redis instances are of the same cloud vendor and in the same region, create a VPC peering connection by referring to VPC Peering Connection.
  - If the VPC of the source and target Redis instances are of the same cloud vendor but in different regions, create a cloud connection by referring to **Cloud Connect Getting Started**.
  - If the source and target Redis instances are on different clouds, create a Direct Connect connection. For details, see **Direct Connect documentation**.

----End

#### **Configuring the Online Migration Task**

- **Step 1** On the **Online Migration** tab page, click **Configure** in the row containing the online migration task you just created.
- **Step 2** Select a migration type.

Supported migration types are **Full** and **Full + Incremental**, which are described in **Table 11-6**.

| Migration Type     | Description  |
|--------------------|--|
| Full               | Suitable for scenarios where services can be interrupted.<br>Data is migrated at one time. <b>Source instance data</b><br><b>updated during the migration will not be migrated to</b><br><b>the target instance.</b>   |
| Full + incremental | Suitable for scenarios requiring minimal service<br>downtime. The incremental migration parses logs to<br>ensure data consistency between the source and target<br>instances.  |
|                    | Once the migration starts, it remains <b>Migrating</b> until you click <b>Stop</b> in the <b>Operation</b> column. After the migration is stopped, data in the source instance will not be lost, but data will not be written to the target instance. When the transmission network is stable, the delay of incremental migration is within seconds. The actual delay depends on the transmission quality of the network link. |

| Table 11- | 6 Migratior | n type descriptior | ۱ |
|-----------|-------------|--------------------|---|
|-----------|-------------|--------------------|---|

#### Figure 11-7 Selecting the migration type

| · Migration Type | • Full<br>Suitable for scenarios where services can be interrupted. Data is migrated at one time. Source Redis data updated during the migration will not be migrated to the target instance. |
|------------------|---|
|                  | Full + Incremental<br>Suitable for scenarios requiring minimal service downtime. The Incremental migration parses logs to ensure data consistency between the source Redis and target Redis.  |

**Step 3** Only if **Migration Type** is set to **Full + Incremental**, you can specify a bandwidth limit.

The data synchronization rate can be kept around the bandwidth limit.

**Step 4** Specify **Auto-Reconnect**. If this option is enabled, automatic reconnections will be performed indefinitely in the case of a network exception.

Full synchronization will be triggered and requires more bandwidth if incremental synchronization becomes unavailable. Exercise caution when enabling this option.

- Step 5 Configure Source Redis and Target Redis.
  - Configure Source Redis Type and Source Redis Instance: Set Redis in the cloud for Source Redis Type and add Source Redis Instance.

If the source Redis is a Redis Cluster, enter the IP addresses and ports of all masters in the cluster and separate multiple addresses with commas (,). For example: **192.168.1.1:6379,192.168.0.0:6379** 

- Configure Target Redis Type and Target Redis Instance:
   Set Redis in the cloud for Target Redis Type and add Target Redis Instance.
- 3. Configure Source Redis Instance Password and Target Redis Instance Password: If the instance is password-protected, click Test Connection to

check whether the instance password is correct and whether the network is connected. If the instance is not password-protected, click **Test Connection** directly. If the test fails, check whether the password is correct, and whether the migration task network is connected.

If a DCS Redis instance is used, the users created in **Managing Users** are currently unavailable.

4. You can specify the source DB and target DB. For example, if you enter **5** for source DB and **6** for target DB, data in DB5 of the source Redis will be migrated to DB6 of the target Redis. If the source DB is not specified but the target DB is specified, all source data will be migrated to the specified target DB by default. If the target DB is not specified, data will be migrated to the corresponding target DB.

#### **NOTE**

- If the source Redis is multi-DB and the target is single-DB (DB0), either ensure that all source data is in DB0, or specify a source DB and set the target DB to 0.
   Otherwise, migration will fail.
- For details about DB in DCS for Redis, see **Does DCS for Redis Support Multi-DB?**

#### Step 6 Click Next.

Step 7 Confirm the migration task details and click Submit.

Go back to the data migration task list. After the migration is successful, the task status changes to **Successful**.

If the migration fails, click the migration task and check the log on the **Migration Logs** page.

D NOTE

- Once incremental migration starts, it remains Migrating after full migration.
- To manually stop a migration task, select the check box on the left of the migration task and click **Stop** above the migration task.
- To perform migration again, select the migration tasks which failed or are stopped, and click **Restart** above. If a restarted migration task fails, click **Configure** to configure the task and try again.
- A maximum of 50 online migration tasks can be selected at a time. You can stop, delete, or restart them in batches.

----End

## Verifying the Migration

Before data migration, if the target Redis has no data, check data integrity after the migration is complete in the following way:

- 1. Connect to the source Redis and the target Redis. Connect to Redis by referring to **redis-cli**.
- 2. Run the **info keyspace** command to check the values of **keys** and **expires**.



3. Calculate the differences between the values of **keys** and **expires** of the source Redis and the target Redis. If the differences are the same, the data is complete and the migration is successful.

During full migration, source Redis data updated during the migration will not be migrated to the target instance.

# 11.4.2 Self-Hosted Redis Migration with Backup Files

This section describes how to migrate self-hosted Redis to DCS by importing backup files.

Simply download the source Redis data and then upload the data to an OBS bucket in the same Huawei Cloud account and region as the target DCS Redis instance. After you have created a migration task on the DCS console, DCS will read data from the OBS bucket and data will be migrated to the target instance.

## Notes and Constraints

- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.
- Migration may fail if the target instance uses smaller specifications than its source.

# Prerequisites

- Before migrating data, read through **Migration Solution Notes** to learn about the DCS data migration function and select an appropriate target instance.
- By default, a Proxy Cluster instance has only one database (DB0). Before you
  migrate data from a multi-DB single-node or master/standby instance to a
  Proxy Cluster instance, check whether any data exists on databases other than
  DB0. If yes, enable multi-DB for the Proxy Cluster instance by referring to
  Enabling Multi-DB.
- By default, a Redis Cluster instance has only one DB (DB0). Before you
  migrate data from a multi-DB single-node or master/standby instance to a
  Redis Cluster instance, check whether any data exists on databases other than
  DB0. To ensure that the migration succeeds, move all data to DB0 by referring
  to Online Migration from Another Cloud Using Rump.
- Prepare the source Redis backup file. The backup file must be in .aof, .rdb, .zip, or .tar.gz format.
- If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.

# **Creating an OBS Bucket and Uploading Backup Files**

If the source Redis backup file to be uploaded is smaller than 5 GB, perform the following steps to create an OBS bucket and upload the file on the OBS console. If

the backup file to be uploaded is larger than 5 GB, upload the file by referring to **instructions**.

**Step 1** Create an OBS bucket on the OBS console.

When creating an OBS bucket, pay attention to the configuration of the following parameters. For details on how to set other parameters, see **Creating a Bucket** in *OBS User Guide*.

1. Region:

The OBS bucket must be in the same region as the target DCS Redis instance.

2. Storage Class: Available options are Standard, Infrequent Access, and Archive.

Do not select Archive. Otherwise, the migration will fail.

- **Step 2** In the bucket list, click the bucket created in **Step 1**.
- Step 3 In the navigation pane, choose Objects.
- Step 4 On the Objects tab page, click Upload Object.
- Step 5 Specify Storage Class.

Do not select Archive. Otherwise, the migration will fail.

Step 6 Upload the objects.

Drag files or folders to the **Upload Object** area or click **add file**.

A maximum of 100 files can be uploaded at a time. The total size cannot exceed 5 GB.

#### Figure 11-8 Uploading an object



- Step 7 Specify Server-Side Encryption. If you enable it, you can select SSE-KMS or SSE-OBS. You can also disable it. For details, see Server-Side Encryption.
- Step 8 Click Upload.

----End

# **Creating a Migration Task**

- **Step 1** Click in the upper left corner and choose **Distributed Cache Service for Redis** under **Middleware** to open the DCS console.
- **Step 2** In the navigation pane, choose **Data Migration**.
- Step 3 Click Create Backup Import Task.
- **Step 4** Enter the task name and description.

The task name must start with a letter, contain 4 to 64 characters, and contain only letters, digits, hyphens (-), and underscores (\_).

- **Step 5** In the **Source Redis** area, select **OBS Bucket** for **Data Source** and then select the OBS bucket to which you have uploaded backup files.
- Step 6 You can specify Source DB to migrate data from the specified DB in the source backup file. For example, if you enter 5, only data in DB5 will be migrated. To migrate all databases, do not specify Source DB.
- **Step 7** Enable **Multi-DB Proxy Cluster** if the source Redis is a multi-DB (**multi-db** set to **yes**) Proxy Cluster DCS Redis instance.
- **Step 8** Click **Add Backup** and select the backup files to be migrated.
- Step 9 In the Target Redis area, select the Target Redis Instance prepared in Prerequisites.
- Step 10 If the target Redis instance has a password, enter the password and click Test Connection to check whether the password is correct. If the instance is not password-protected, click Test Connection directly.
- Step 11 For Target DB, you can specify a DB in the target Redis to migrate data to. For example, if you enter 5, data will be migrated to DB5 of the target Redis. If you do not specify a DB, data will be migrated to a DB corresponding to the source DB.

- If the source Redis is multi-DB and the target is single-DB (DB0), either ensure that all source data is in DB0, or specify a source DB and set the target DB to **0**. Otherwise, migration will fail.
- For details about DB in DCS for Redis, see Does DCS for Redis Support Multi-DB?
- Step 12 Click Next.
- **Step 13** Confirm the migration task details and click **Submit**.

Go back to the data migration task list. After the migration is successful, the task status changes to **Successful**.

----End

# Verifying the Migration

After the data is imported successfully, access the DCS instance and run the **info** command to check whether the data has been successfully imported as required. Connect to Redis by referring to **redis-cli**.

If the import fails, check the procedure. If the import command is incorrect, run the **flushall** or **flushdb** command to clear the cache data in the target instance, modify the import command, and try again.

# 11.4.3 Self-Hosted Redis Migration with redis-cli (AOF)

redis-cli is the command line tool of Redis, which can be used after you install the Redis server. This section describes how to use redis-cli to migrate a data from a self-hosted Redis instance to a DCS instance.

An AOF file can be generated quickly. It applies to scenarios where you can access the Redis server and modify the configurations, such as scenarios with self-built Redis servers.

# Notes and Constraints

- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.
- Migration may fail if the target instance uses smaller specifications than its source.
- Migrate data during off-peak hours.
- Before data migration, suspend your services so that newly generated data changes will not be lost during the migration.

## Prerequisites

- If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.
- An Elastic Cloud Server (ECS) has been created. For details about how to create an ECS, see **Purchasing a Custom ECS**.

# Generating an AOF File

- 1. Log in to the ECS.
- 2. Install redis-cli. The following steps assume that your client is installed on the Linux OS.
  - a. Run the following command to download Redis: You can also install other Redis versions. For details, see the **Redis official website**. wget http://download.redis.io/releases/redis-5.0.8.tar.gz
  - b. Run the following command to decompress the source code package of your Redis client:

tar -xzf redis-5.0.8.tar.gz

- c. Run the following commands to go to the Redis directory and compile the source code of your Redis client: cd redis-5.0.8 make cd src
- 3. Run the following command to enable cache persistence and obtain the AOF persistence file:

. redis-cli -h {*source\_redis\_address*} -p {*port*} -a {*password*} config set appendonly yes

*{source\_redis\_address}* is the connection address of the source Redis, *{port}* is the port of the source Redis, and *{password}* is the connection password of the source Redis.

If the size of the AOF file does not change after you have enabled persistence, the AOF file contains full cached data.

#### **NOTE**

- To find out the path for storing the AOF file, use redis-cli to access the Redis instance, and run the **config get dir** command. Unless otherwise specified, the file is named as **appendonly.aof** by default.
- To disable synchronization after the AOF file is generated, use redis-cli to log in to the Redis instance and run the **config set appendonly no** command.

# Uploading the AOF file to Huawei Cloud ECS

To save time, you are advised to compress the AOF file and upload it to Huawei Cloud ECS using an appropriate mode (for example, SFTP mode).

#### **NOTE**

Ensure that the ECS has sufficient disk space for data file decompression, and can communicate with the DCS instance. Generally, the ECS and DCS instance are configured to belong to the same VPC and subnet, and the configured security group rules do not restrict access ports. For details about how to configure a security group, see **Security Group Configurations**.

#### **Importing Data**

Log in to the ECS and run the following command to import data.

redis-cli -h {*dcs\_instance\_address*} -p {*port*} -a {*password*} --pipe < appendonly.aof

*{dcs\_instance\_address}* indicates the address of the target Redis instance, *{port}* indicates the port of the target Redis instance, and *{password}* indicates the password for connecting to the target Redis instance.

It takes 4 to 10 seconds to import an AOF file of 1 million data (20 bytes per data segment) to a VPC.

#### NOTICE

If SSL is enabled, replace the instance address and port number with the actual values.
# Verifying the Migration

After the data is imported successfully, access the DCS instance and run the **info** command to check whether the data has been successfully imported as required. Connect to Redis by referring to **redis-cli**.

If the import fails, check the procedure. If the import command is incorrect, run the **flushall** or **flushdb** command to clear the cache data in the target instance, modify the import command, and try again.

# 11.4.4 Self-Hosted Redis Migration with redis-cli (RDB)

redis-cli is the command line tool of Redis, which can be used after you install the Redis server. redis-cli supports data export as an RDB file. If your Redis service does not support AOF file export, use redis-cli to obtain an RDB file. Then, use another tool (such as redis-shake) to import the file to a DCS instance.

# Notes and Constraints

- Migrate data during off-peak hours.
- When the source is Redis native cluster data, individually export the data of each node in the cluster, and then import the data node by node.
- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

# Prerequisites

- If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.
- An Elastic Cloud Server (ECS) has been created. For details about how to create an ECS, see **Purchasing a Custom ECS**.
- The source self-hosted Redis instance must support the SYNC command. Otherwise, the RDB file cannot be exported using redis-cli.

# **Exporting the RDB File**

1. Prepare for the export.

For master/standby or cluster DCS instances, there is a delay in writing data into an RDB file based on the delay policies configured in the **redis.conf** file. Before data export, learn the RDB policy configurations of the Redis instance to be migrated, suspend your service systems, and then write the required number of test data into the Redis instance. This ensures that the RDB file is newly generated.

For example, the default RDB policy configurations in the **redis.conf** file are as follows:

save 900 1 //Writes changed data into an RDB file if there is any data change within 900s. save 300 10 //Writes changed data into an RDB file if there are more than 10 data changes within

300s.

save 60 10000 //Writes changed data into an RDB file if there are more than 10,000 data changes within 60s.

Based on the preceding policy configurations, after stopping your service systems from writing data into the Redis instances, you can write test data to trigger the policies, so that all service data can be synchronized to the RDB file.

You can delete the test data after data import.

**NOTE** 

- If there is any DB not used by your service systems, you can write test data into the DB, and run the **flushdb** command to clear the database after importing data into DCS.
- Compared with master/standby instances, single-node instances without data persistence configured require a longer time for export of an RDB file, because the RDB file is temporarily generated.
- 2. Log in to the ECS.
- 3. Install redis-cli. The following steps assume that your client is installed on the Linux OS.
  - a. Run the following command to download Redis: You can also install other Redis versions. For details, see the **Redis official website**. wget http://download.redis.io/releases/redis-5.0.8.tar.gz
  - Run the following command to decompress the source code package of your Redis client: tar -xzf redis-5.0.8.tar.gz
  - c. Run the following commands to go to the Redis directory and compile the source code of your Redis client: cd redis-5.0.8 make cd src
- 4. Run the following command to export the RDB file: redis-cli -h {source\_redis\_address} -p {port} -a {password} --rdb {output.rdb}

*{source\_redis\_address}* is the connection address of the source Redis, *{port}* is the port of the source Redis, *{password}* is the connection password of the source Redis, and *{output.rdb}* is the RDB file name.

If "Transfer finished with success." is displayed after the command is executed, the file is exported successfully.

# Uploading the RDB File to Huawei Cloud ECS

To save time, you are advised to compress the RDB file and upload it to Huawei Cloud ECS using an appropriate mode (for example, SFTP mode).

**NOTE** 

Ensure that the ECS has sufficient disk space for data file decompression, and can communicate with the DCS instance. Generally, the ECS and DCS instance are configured to belong to the same VPC and subnet, and the configured security group rules do not restrict access ports. For details about how to configure a security group, see **Security Group Configurations**.

# Importing Data

Use redis-shake to import data.

It takes 4 to 10 seconds to import an RDB file of 1 million data (20 bytes per data segment) to a VPC.

# Verifying the Migration

After the data is imported successfully, access the DCS instance and run the **info** command to check whether the data has been successfully imported as required. Connect to Redis by referring to **redis-cli**.

If the import fails, check the procedure. If the import command is incorrect, run the **flushall** or **flushdb** command to clear the cache data in the target instance, modify the import command, and try again.

# 11.4.5 Self-Hosted Redis Cluster Migration with redis-shake (Online)

redis-shake is an open-source tool for migrating data online or offline (by importing backup files) between Redis Clusters. Data can be migrated to DCS Redis Cluster instances seamlessly because DCS Redis Cluster inherits the native Redis Cluster design.

The following describes how to use Linux redis-shake to migrate self-hosted Redis Cluster to a DCS Redis Cluster instance online.

# **Notes and Constraints**

- To migrate data from a self-hosted Redis Cluster instance to a DCS Redis Cluster instance online, ensure that the source Redis is connected to the target Redis, or use a transit cloud server to connect the source and target cluster instances.
- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

# Prerequisites

- If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.
- An Elastic Cloud Server (ECS) has been created. For details about how to create an ECS, see **Purchasing a Custom ECS**.

Select the same VPC, subnet, and security group as the DCS Redis Cluster instance, and bind EIPs to the ECS.

• If the source self-hosted Redis Cluster is deployed on cloud servers of another cloud, allow public access to the servers.

# **Obtaining Information of the Source and Target Redis Nodes**

- 1. Connect to the source and target Redis instances, respectively. Connect to Redis by referring to **redis-cli**.
- 2. In online migration of Redis Clusters, the migration must be performed node by node. Run the following command to query the IP addresses and ports of all nodes in both the source and target Redis Clusters. redis-cli -h {redis\_address} -p {redis\_port} -a {redis\_password} cluster nodes

*{redis\_address}* indicates the Redis connection address, *{redis\_port}* indicates the Redis port, and *{redis\_password}* indicates the Redis connection password.

In the output similar to the following, obtain the IP addresses and port numbers of all the master nodes.



# Configuring the redis-shake Tool

- 1. Log in to the ECS.
- 2. Run the following command on the ECS to download the redis-shake: This section uses v4.3.2 as an example. You can also download **other redis-shake versions** as required.

wget https://github.com/tair-opensource/RedisShake/releases/download/v4.3.2/redis-shake-v4.3.2-linux-amd64.tar.gz

3. Decompress the redis-shake file. mkdir redis-shake-v4.3.2 tar -C redis-shake-v4.3.2 -xzvf redis-shake-v4.3.2-linux-amd64.tar.gz

| [root@ecs    |          |        | redis-sha | ake-N | /4.3 | 3.2]# 1 | 11          |
|--------------|----------|--------|-----------|-------|------|---------|-------------|
| total 11516  |          |        |           |       |      |         |             |
| -rwxr-xr-x 1 | sysadmin | docker | 11783865  | Jan   | 14   | 19:04   | redis-shake |
| -rw-rr 1     | sysadmin | docker | 6696      | Jan   | 14   | 19:04   | shake.toml  |

- 4. Go to the decompressed directory. cd redis-shake-v4.3.2
- 5. Edit the **shake.toml** file by providing the following information of both the source and the target. vim shake.toml

The modification is as follows: [sync\_reader] # If the source instance type is a Redis Cluster, set the value to **true**. cluster = true # IP address and port of any node in the source Redis Cluster address = {redis\_ip}:{redis\_port} # If there is no password, skip the following parameter password = {source\_redis\_password} [redis\_writer] # If the target instance type is a Redis Cluster, set the value to **true**. cluster = true # IP address and port of any node in the target Redis Cluster address = {redis\_ip}:{redis\_port} # If there is no password, skip the following parameter password = {target\_redis\_password}

Press **Esc** to exit the editing mode and enter **:wq!**. Press **Enter** to save the configuration and exit the editing interface.

# Migrating Data Online

Run the following command to synchronize data between the source and the target Redis: ./redis-shake shake.toml

If the following information is displayed, the full synchronization has been completed and incremental synchronization begins.

syncing aof

If the following information is displayed, no new data is incremented. You can stop the incremental synchronization by pressing **Ctrl+C**.

write\_ops=[0.00], src-\*, syncing aof, diff=[0]

Figure 11-9 Online migration using redis-shake

| [root@ecs-xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx  |
|---|
| 2025-03-11 12:30:07 INF load config from file: shake.toml   |
| 2025-03-11 12:30:07 INF log level: [info], log file: [/tmp/redis-shake-v4.3.2/data/shake.log]   |
| 2025-03-11 12:30:07 INF changed work dir. dir=[/tmp/redis-shake-v4.3.2/data]  |
| 2025-03-11 12:30:07 INF GOMAXPROCS defaults to the value of runtime.NumCPU [2]  |
| 2025-03-11 12:30:07 INF not set pprof port  |
| 2025-03-11 12:30:07 INF create SyncClusterReader  |
| 2025-03-11 12:30:07 INF * address (should be the address of one node in the Redis cluster): 127.0.0.1:8715  |
| 2025-03-11 12:30:07 INF * username:   |
| 2025-03-11 12:30:07 INF * password:   |
| 2025-03-11 12:30:07 INF * tls: false  |
| 2025-03-11 12:30:07 INF address=127.0.0.1:8715, reply=6c9438173d174f3daea3b5f49f795421177b64ad 127.0.0.1:8785@18785 slave 2d7fcb6006ffc2e1f467f9ac0920d40c92fe8bf5 0 1741667406205 9 connecte |
| 07e965756a364d856756f349bb89f23ac956a2ac 127.0.0.1:8775@18775 slave 853d0b57a2ea9fafa5025e41f13fd453e62b4fe5 0 1741667404000 8 connected  |
| 22cf6d0718b58fd814c4c858259eb0a6f8b132a4 127.0.0.1:8745@18745 slave 63cfa92b7506cd7bd54972ba73259b1d4d617825 0 1741667403197 10 connected   |
| 26be57d088e63088acf50b4b1b98dc10bf0d4ada 127.0.0.1:8715@18715 myself,slave 63cfa92b7506cd7bd54972ba73259b1d4d617825 0 1741667401000 2 connected   |
| 63cfa92b7506cd7bd54972ba73259b1d4d617825 127.0.0.1:8755@18755 master - 0 1741667404000 10 connected 5461-10922  |
| 932761fae199672c152e32b64d0f564b6e1fe878 127.0.0.1:8765@18765 slave 2d7fcb6006ffc2e1f467f9ac0920d40c92fe8bf5 0 1741667405202 7 connected  |
| 853d0b57a2ea9fafa5025e41f13fd453e62b4fe5 127.0.0.1:8705@18705 master - 0 1741667404199 1 connected 0-5460   |
| 15f97334e9aaa0f92b08e1ee6d0f82a30907d989 127.0.0.1:8735@18735 slave 853d0b57a2ea9fafa5025e41f13fd453e62b4fe5 0 1741667403000 4 connected  |
| 2d7fcb6006ffc2e1f467f9ac0920d40c92fe8bf5 127.0.0.1:8725@18725 master - 0 1741667405000 3 connected 10923-16383  |
| 2025-03-11 12:30:07 INF create RedisClusterWriter   |
| 2025-03-11 12:30:07 INF * address (should be the address of one node in the Redis cluster): 127.0.0.1:8716  |
| 2025-03-11 12:30:07 INF * username:   |
| 2025-03-11 12:30:07 INF * password:   |
| 2025-03-11 12:30:07 INF * tls: false  |
| 2025-03-11 12:30:07 INF address=127.0.0.1:8716, reply=68090e42e3242da51ab1af8327cf6a95d3f971ed 127.0.0.1:8726@18726 master - 0 1741667405080 2 connected 5461-10922                           |
| 07eb7c2d21029ee9986601476b1b77caee8a8bf6 127.0.0.1:8736@18736 master - 0 1741667406083 3 connected 10923-16383  |
| 142ac36f1015bbf2adaf29c676b614cdef6f07e5 127.0.0.1:8756@18756 slave 07eb7c2d21029ee9986601476b1b77caee8a8bf6 0 1741667402070 3 connected  |
| 5ba9fb74da45602186b7b0a67ccc5d776d162362 127.0.0.1:8746@18746 slave 68090e42e3242da51ab1af8327cf6a95d3f971ed 0 1741667403074 2 connected  |
| 77b9cabb8db8cab17b93392a437e7d1f909a33da 127.0.0.1:8766@18766 slave 2376279f28bb83946d05fcf9b87638d6f4f6f989 0 1741667404076 1 connected  |
| 2376279f28bb83946d05fcf9b87638d6f4f6f989 127.0.0.1:8716@18716 myself,master - 0 1741667404000 1 connected 0-5460  |
| 2025-03-11 12:30:07 INF redisClusterWriter connected to redis cluster successful. addresses=[127.0.0.1:8726 127.0.0.1:8736 127.0.0.1:8716]  |
| 2025-03-11 12:30:07 INF start syncing   |
| 2025-03-11 12:30:07 INF [reader_127.0.0.1_8755] source db is not doing bgsave! continue.  |
| 2025-03-11 12:30:07 INF [reader_127.0.0.1_8725] source db is not doing bgsave! continue.  |
| 2025-03-11 12:30:07 INF [reader_127.0.0.1_8705] source db is not doing bgsave! continue.  |
| 2025-03-11 12:30:12 INF read_count=[2304], read_ops=[0.00], write_count=[2304], write_ops=[0.00], src-1, syncing aof, diff=[944118872]  |
| 2025-03-11 12:30:17 INF read_count=[2304], read_ops=[460.82], write_count=[2304], write_ops=[460.82], src-2, syncing aof, diff=[0]  |
| 2025-03-11 12:30:22 INF read_count=[2304], read_ops=[0.00], write_count=[2304], write_ops=[0.00], src-0, syncing aof, diff=[0]  |
| 2025-03-11 12:30:27 INF read_count=[2304], read_ops=[0.00], write_count=[2304], write_ops=[0.00], src-1, syncing aof, diff=[0]  |
| 2025-03-11 12:30:32 INF read_count=[2304], read_ops=[0.00], write_count=[2304], write_ops=[0.00], src-2, syncing aof, diff=[0]  |
| 2025-03-11 12:30:37 INF read count=[2304], read ops=[0.00], write count=[2304], write ops=[0.00], src-0, syncing aof, diff=[0]  |

# Verifying the Migration

- 1. After the data synchronization, connect to the Redis Cluster DCS instance by referring to **redis-cli**.
- 2. Run the **info** command to check whether the data has been successfully imported as required.

If the data has not been fully imported, run the **flushall** or **flushdb** command to clear the cached data in the target instance, and migrate data again.

3. After the verification is complete, you are advised to clear the redis-shake configuration in time.

# 11.4.6 Self-Hosted Redis Cluster Migration with redis-shake (RDB)

redis-shake is an open-source tool for migrating data online or offline (by importing backup files) between Redis Clusters. Data can be migrated to DCS Redis Cluster instances seamlessly because DCS Redis Cluster inherits the native Redis Cluster design.

The following describes how to use Linux redis-shake to migrate self-hosted Redis Cluster to a DCS Redis Cluster instance offline.

# D NOTE

If the source Redis and the target Redis cannot be connected, or the source Redis is deployed on other clouds, you can migrate data by importing backup files.

# Notes and Constraints

To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

# Prerequisites

• A DCS Redis Cluster instance has been created. For details about how to create one, see **Buying a DCS Redis Instance**.

The memory of the target Redis instance cannot be smaller than that of the source Redis.

• An Elastic Cloud Server (ECS) has been created. For details about how to create an ECS, see **Purchasing an ECS**. Select the same VPC, subnet, and security group as the DCS Redis Cluster instance.

# **Obtaining Information of the Source and Target Redis Nodes**

- 1. Connect to the source and target Redis instances, respectively. Connect to Redis by referring to **redis-cli**.
- 2. In online migration of Redis Clusters, the migration must be performed node by node. Run the following command to query the IP addresses and ports of all nodes in both the source and target Redis Clusters. redis-cli -h {redis\_address} -p {redis\_port} -a {redis\_password} cluster nodes

*{redis\_address}* indicates the Redis connection address, *{redis\_port}* indicates the Redis port, and *{redis\_password}* indicates the Redis connection password.

In the output similar to the following, obtain the IP addresses and port numbers of all the master nodes.

| [root@ecs-         | 54-centos ~]# redi      | s-cli -h 192.1 | 68.0.140 -p 6379  | -a <b>@</b> 23 ( | luster nodes               |
|--------------------|-------------------------|----------------|-------------------|------------------|----------------------------|
| fb75f0743af4695a3d | l241ff7790b2f508e4985ff | 192.168.0.140  | :6379@16379 myse  | lf,master - 0    | 1562144170000 3 connected  |
| d112bae791b2bbd960 | 2fe32963536b8a0db9eb79  | 192.168.0.61:  | 6379@16379 maste  | r - 0 1562144:   | 171524 1 connected 0-5460  |
| 73e2f8fe196166f9ad | 1283361867d24c136413f0  | 192.168.0.194  | :6379@16379 mast  | er - 0 1562144   | 170000 2 connected 5461-1  |
| 40d72299fde6045de0 | f79ee4b97910b505acbc6a  | 192.168.0.231  | :6379@16379 slave | e 73e2f8fe196:   | L66f9ad1283361867d24c13641 |
| be6c07faa64d724323 | e0d7cedc3f38346dcbd212  | 192.168.0.80:  | 6379@16379 slave  | fb75f0743af4     | 595a3d241ff7790b2f508e4985 |
| c16b9acaeed7dd0721 | f129596cd43bd499c0e396  | 192.168.0.169  | :6379@16379 slav  | e d112bae791b    | 2bbd9602fe32963536b8a0db9e |

# Installing RedisShake

- 1. Log in to the ECS.
- 2. Run the following command on the ECS to download the redis-shake: This section uses v2.1.2 as an example. You can also download **other redis-shake versions** as required.

wget https://github.com/tair-opensource/RedisShake/releases/download/release-v2.1.2-20220329/ release-v2.1.2-20220329.tar.gz

3. Decompress the redis-shake file. tar -xvf release-v2.1.2-20220329.tar.gz

| [root@ecs-437a tem]# tar -xvf release-v2.1.2-20220329.tar.gz |
|--|
| bin/redis-shake.conf   |
| bin/redis-shake.darwin                                       |
| bin/redis-shake.linux  |
| bin/redis-shake.windows                                      |
| [root@ecs-437a tem]# ll                                      |
| total 17960  |
| drwxr-xr-x 2 root root     4096 Apr 22 19:28 bin             |
| -rw-rr 1 root root 18383025                                  |

### D NOTE

If the source cluster is deployed in the data center intranet, install redis-shake on the intranet server. Export the source cluster backup file by referring to **Exporting the Backup File**. Upload the backup to the cloud server as instructed by the following steps

# **Exporting the Backup File**

1. Go to the redis-shake directory.

cd bin

| [root@ecs-437a tem]# cd bin |   |     |       |          |     |    |      |                     |
|-----------------------------|---|-----|-------|----------|-----|----|------|---------------------|
| [root@ecs-437a bin]# ll     |   |     |       |          |     |    |      |                     |
| total 34776                 | 5 |     |       |          |     |    |      |                     |
| -rw-rr                      | 1 | 502 | games | 13693    | Mar | 29 | 2022 | redis-shake.conf    |
| -rwxr-xr-x                  | 1 | 502 | games | 11740624 | Mar | 29 | 2022 | redis-shake.darwin  |
| -rwxr-xr-x                  | 1 | 502 | games | 11797281 | Mar | 29 | 2022 | redis-shake.linux   |
| -rwxr-xr-x                  | 1 | 502 | games | 12048384 | Mar | 29 | 2022 | redis-shake.windows |

2. Edit the **redis-shake.conf** file by providing the following information about all the masters of the source: vim redis-shake.conf

The modification is as follows: source.type = cluster # If there is no password, skip the following parameter. source.password\_raw = {source\_redis\_password} # IP addresses and port numbers of all masters of the source Redis Cluster, which are separated by semicolons (;). source.address = {master1\_ip}:{master1\_port};{master2\_ip}:{master2\_port}...{masterN\_ip}: {masterN\_port}

Press **Esc** to exit the editing mode and enter **:wq!**. Press **Enter** to save the configuration and exit the editing interface.

3. Run the following command to export the RDB file: ./redis-shake -type dump -conf redis-shake.conf

If the following information is displayed in the execution log, the backup file is exported successfully:

execute runner[\*run.CmdDump] finished!

# Importing the Backup File

- 1. Import the RDB file (or files) to the cloud server. The cloud server must be connected to the target DCS instance.
- Edit the redis-shake.conf file by providing the following information about all the masters of the target: vim redis-shake.conf

```
The modification is as follows:

target.type = cluster

# If there is no password, skip the following parameter.

target.password_raw = {target_redis_password}

# IP addresses and port numbers of all masters of the target instance, which are separated by

semicolons (;).

target.address = {master1_ip}:{master1_port};{master2_ip}:{master2_port}...{masterN_ip}:

{masterN_port}

# List the RDB files to be imported, separated by semicolons (;).

rdb.input = {local_dump.0};{local_dump.1};{local_dump.2};{local_dump.3}
```

Press **Esc** to exit the editing mode and enter **:wq!**. Press **Enter** to save the configuration and exit the editing interface.

3. Run the following command to import the RDB file to the target instance: ./redis-shake -type restore -conf redis-shake.conf

If the following information is displayed in the execution log, the backup file is imported successfully:

Enabled http stats, set status (incr), and wait forever.

# Verifying the Migration

- 1. After the data synchronization, connect to the Redis Cluster DCS instance by referring to **redis-cli**.
- 2. Run the **info** command to check whether the data has been successfully imported as required.

If the data has not been fully imported, run the **flushall** or **flushdb** command to clear the cached data in the target instance, and migrate data again.

3. After the verification is complete, you are advised to clear the redis-shake configuration in time.

# **11.5 Migration from Another Cloud**

# 11.5.1 Migrating Redis from Another Cloud Online

If the source and target instances are interconnected and the **SYNC** and **PSYNC** commands are supported by the source instance, data can be migrated online in full or incrementally from another cloud to the target DCS.

# **NOTE**

During online migration, data is essentially synchronized in full to a new replica. Therefore, perform online migration during low-demand hours. Otherwise, source instance CPU usage may surge and latency may increase.

# **Notes and Constraints**

- If the **SYNC** and **PSYNC** commands are disabled by the source instance, enable them by contacting the source vendor. Otherwise, the migration fails.
- You cannot use public networks for online migration.
- During online migration, you are advised to set **repl-timeout** on the source instance to 300s and **client-output-buffer-slave-hard-limit** and **client-output-buffer-slave-soft-limit** to 20% of the maximum memory of the source instance.
- The source must be Redis 3.0 or later.
- For earlier instances whose passwords contain single quotation marks ('), modify the password for online migration or try other methods.
- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

# Prerequisites

• Before migrating data, read through **Migration Solution Notes** to learn about the DCS data migration function and select an appropriate target instance.

- By default, a Proxy Cluster instance has only one database (DB0). Before you
  migrate data from a multi-DB single-node or master/standby instance to a
  Proxy Cluster instance, check whether any data exists on databases other than
  DB0. If yes, enable multi-DB for the Proxy Cluster instance by referring to
  Enabling Multi-DB.
- By default, a Redis Cluster instance has only one DB (DB0). Before you
  migrate data from a multi-DB single-node or master/standby instance to a
  Redis Cluster instance, check whether any data exists on databases other than
  DB0. To ensure that the migration succeeds, move all data to DB0 by referring
  to Online Migration from Another Cloud Using Rump.
- The IP address and port of the source Redis instance has been obtained.
- If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.

# **Creating an Online Migration Task**

- **Step 1** Log in to the DCS console using the account of the target DCS Redis instance.
- **Step 2** Click <sup>Q</sup> in the upper left corner of the console and select the region where your target instance is located.
- **Step 3** In the navigation pane, choose **Data Migration**.
- **Step 4** Click **Create Online Migration Task**.
- **Step 5** Enter the task name and description.

The task name must start with a letter, contain 4 to 64 characters, and contain only letters, digits, hyphens (-), and underscores (\_).

**Step 6** Configure the VPC, subnet, and security group for the migration task.

**NOTE** 

- Select the same VPC as the target Redis. Ensure that the migration resource can access the target Redis instance.
- The online migration task uses a tenant IP address (**Migration ECS** displayed on the **Basic Information** page of the task.) If a whitelist is configured for the source or target instance, add the migration IP address to the whitelist or disable the whitelist.
- To allow the VM used by the migration task to access the source and target instances, set an outbound rule for the task's security group to allow traffic through the IP addresses and ports of the source and target instances. By default, all outbound traffic is allowed.

----End

# **Checking the Network**

**Step 1** Check whether the source Redis instance, the target Redis instance, and the migration task are configured with the same VPC.

If yes, go to **Configuring the Online Migration Task**. If no, go to **Step 2**.

**Step 2** Check whether the VPCs configured for the source Redis instance, the target Redis instance, and the migration task are connected to ensure that the VM resource of the migration task can access the source and target Redis instances.

If yes, go to **Configuring the Online Migration Task**. If no, go to **Step 3**.

**Step 3** If the source and target Redis instances are on different clouds, create a Direct Connect connection. For details, see **Direct Connect documentation**.

----End

\* Miorati

# **Configuring the Online Migration Task**

- **Step 1** On the **Online Migration** tab page, click **Configure** in the row containing the online migration task you just created.
- **Step 2** Select a migration type.

Supported migration types are **Full** and **Full + Incremental**, which are described in **Table 11-7**.

| Migration Type     | Description  |
|--------------------|--|
| Full               | Suitable for scenarios where services can be interrupted.<br>Data is migrated at one time. <b>Source instance data</b><br><b>updated during the migration will not be migrated to</b><br><b>the target instance.</b>   |
| Full + incremental | Suitable for scenarios requiring minimal service<br>downtime. The incremental migration parses logs to<br>ensure data consistency between the source and target<br>instances.  |
|                    | Once the migration starts, it remains <b>Migrating</b> until you click <b>Stop</b> in the <b>Operation</b> column. After the migration is stopped, data in the source instance will not be lost, but data will not be written to the target instance. When the transmission network is stable, the delay of incremental migration is within seconds. The actual delay depends on the transmission quality of the network link. |

|  | Table | 11-7 | Migration | type | description |
|--|-------|------|-----------|------|-------------|
|--|-------|------|-----------|------|-------------|

#### Figure 11-10 Selecting the migration type

| on Type | • Full<br>Suitable for scenarios where services can be interrupted. Data is migrated at one time. Source Redis data updated during the migration will not be migrated to the target instance. |
|---------|---|
|         | Full + Incremental<br>Suitable for scenarios requiring minimal service downtime. The Incremental migration parses logs to ensure data consistency between the source Redis and target Redis.  |

**Step 3** Only if **Migration Type** is set to **Full + Incremental**, you can specify a bandwidth limit.

The data synchronization rate can be kept around the bandwidth limit.

**Step 4** Specify **Auto-Reconnect**. If this option is enabled, automatic reconnections will be performed indefinitely in the case of a network exception.

Full synchronization will be triggered and requires more bandwidth if incremental synchronization becomes unavailable. Exercise caution when enabling this option.

#### Step 5 Configure Source Redis and Target Redis.

1. Configure Source Redis Type and Source Redis Instance:

Set **Redis in the cloud** for **Source Redis Type** and add **Source Redis Instance**.

If the source Redis is a Redis Cluster, enter the IP addresses and ports of all masters in the cluster and separate multiple addresses with commas (,). For example: **192.168.1.1:6379,192.168.0.0:6379** 

2. Configure Target Redis Type and Target Redis Instance:

Set Redis in the cloud for Target Redis Type and add Target Redis Instance.

3. Configure **Source Redis Instance Password** and **Target Redis Instance Password**: If the instance is password-protected, click **Test Connection** to check whether the instance password is correct and whether the network is connected. If the instance is not password-protected, click **Test Connection** directly. If the test fails, check whether the password is correct, and whether the migration task network is connected.

If a DCS Redis instance is used, the users created in **Managing Users** are currently unavailable.

4. You can specify the source DB and target DB. For example, if you enter **5** for source DB and **6** for target DB, data in DB5 of the source Redis will be migrated to DB6 of the target Redis. If the source DB is not specified but the target DB is specified, all source data will be migrated to the specified target DB by default. If the target DB is not specified, data will be migrated to the corresponding target DB.

#### **NOTE**

- If the source Redis is multi-DB and the target is single-DB (DB0), either ensure that all source data is in DB0, or specify a source DB and set the target DB to 0.
   Otherwise, migration will fail.
- For details about DB in DCS for Redis, see **Does DCS for Redis Support Multi-DB?**

#### Step 6 Click Next.

**Step 7** Confirm the migration task details and click **Submit**.

Go back to the data migration task list. After the migration is successful, the task status changes to **Successful**.

If the migration fails, click the migration task and check the log on the **Migration Logs** page.

# 

- Once incremental migration starts, it remains **Migrating** after full migration.
- To manually stop a migration task, select the check box on the left of the migration task and click **Stop** above the migration task.
- To perform migration again, select the migration tasks which failed or are stopped, and click **Restart** above. If a restarted migration task fails, click **Configure** to configure the task and try again.
- A maximum of 50 online migration tasks can be selected at a time. You can stop, delete, or restart them in batches.

#### ----End

# Verifying the Migration

Before data migration, if the target Redis has no data, check data integrity after the migration is complete in the following way:

- 1. Connect to the source Redis and the target Redis. Connect to Redis by referring to **redis-cli**.
- 2. Run the **info keyspace** command to check the values of **keys** and **expires**.

|   | 192.1                              |
|---|------------------------------------|
|   | # Keyspace                         |
|   | db0:keys=81869,expires=0,avg_ttl=0 |
| н | 192.14.217:6379>                   |

3. Calculate the differences between the values of **keys** and **expires** of the source Redis and the target Redis. If the differences are the same, the data is complete and the migration is successful.

During full migration, source Redis data updated during the migration will not be migrated to the target instance.

# 11.5.2 Backup Import from Another Cloud

This section describes how to migrate Redis from another cloud to DCS by importing backup files.

Simply download the source Redis data and then upload the data to an OBS bucket in the same account and region as the target DCS Redis instance. After you have created a migration task on the DCS console, DCS will read data from the OBS bucket and data will be migrated to the target instance.

# Notes and Constraints

- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see Transmitting DCS Redis Data with Encryption Using SSL.
- Migration may fail if the target instance uses smaller specifications than its source.

# Prerequisites

• Before migrating data, read through **Migration Solution Notes** to learn about the DCS data migration function and select an appropriate target instance.

- By default, a Proxy Cluster instance has only one database (DB0). Before you
  migrate data from a multi-DB single-node or master/standby instance to a
  Proxy Cluster instance, check whether any data exists on databases other than
  DB0. If yes, enable multi-DB for the Proxy Cluster instance by referring to
  Enabling Multi-DB.
- By default, a Redis Cluster instance has only one DB (DB0). Before you migrate data from a multi-DB single-node or master/standby instance to a Redis Cluster instance, check whether any data exists on databases other than DB0. To ensure that the migration succeeds, move all data to DB0 by referring to Online Migration from Another Cloud Using Rump.
- Prepare the source Redis backup file. The backup file must be in .aof, .rdb, .zip, or .tar.gz format.
- If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.

# Creating an OBS Bucket and Uploading Backup Files

If the source Redis backup file to be uploaded is smaller than 5 GB, perform the following steps to create an OBS bucket and upload the file on the OBS console. If the backup file to be uploaded is larger than 5 GB, upload the file by referring to **instructions**.

**Step 1** Create an OBS bucket on the OBS console.

When creating an OBS bucket, pay attention to the configuration of the following parameters. For details on how to set other parameters, see **Creating a Bucket** in *OBS User Guide*.

1. Region:

The OBS bucket must be in the same region as the target DCS Redis instance.

2. **Storage Class**: Available options are **Standard**, **Infrequent Access**, and **Archive**.

Do not select Archive. Otherwise, the migration will fail.

- **Step 2** In the bucket list, click the bucket created in **Step 1**.
- Step 3 In the navigation pane, choose Objects.
- **Step 4** On the **Objects** tab page, click **Upload Object**.
- Step 5 Specify Storage Class.

Do not select **Archive**. Otherwise, the migration will fail.

Step 6 Upload the objects.

Drag files or folders to the **Upload Object** area or click **add file**.

A maximum of 100 files can be uploaded at a time. The total size cannot exceed 5 GB.

| opioad object          | (Optional) Co                             | nfigure Advanced Settings  |  |   |
|------------------------|---|--|--|---|
| Storage Class          | Standard                                  | Infrequent Access  | Archive  |   |
|                        | Optimized for freque                      | ntly accessed (multiple times                                    | per month) data such as small  | and essential files that require low latency. |
|                        | If you do not change bucket creation. Lea | this setting, your uploaded ob<br>rn more                        | jects will be stored using the d                                     | efault storage class you selected during      |
| Jpload Object          | A The file or fol versions of the         | der you newly upload will ove<br>e same file or folder, enable v | write any existing file or folder<br>ersioning for the current bucke | with the same name. To keep different<br>t.   |
|                        |   |  | OBS  |   |
|                        |   | D  |  | _   |
|                        | (A  | Drag and<br>maximum of 100 files can be                          | drop files or folders, or add fil<br>uploaded at a time. The total s | es<br>ze cannot exceed 5 GB.)                 |
| Server-Side Encryption | (A<br>SSE-KMS                             | Drag and<br>maximum of 100 files can be<br>SSE-OBS               | drop files or folders, or add fil<br>uploaded at a time. The total s | es<br>ize cannot exceed 5 GB.)                |

Figure 11-11 Uploading an object

- **Step 7** Specify **Server-Side Encryption**. If you enable it, you can select **SSE-KMS** or **SSE-OBS**. You can also disable it. For details, see **Server-Side Encryption**.
- Step 8 Click Upload.

----End

# **Creating a Migration Task**

- **Step 1** Click in the upper left corner and choose **Distributed Cache Service for Redis** under **Middleware** to open the DCS console.
- **Step 2** In the navigation pane, choose **Data Migration**.
- Step 3 Click Create Backup Import Task.
- **Step 4** Enter the task name and description.

The task name must start with a letter, contain 4 to 64 characters, and contain only letters, digits, hyphens (-), and underscores (\_).

- **Step 5** In the **Source Redis** area, select **OBS Bucket** for **Data Source** and then select the OBS bucket to which you have uploaded backup files.
- Step 6 You can specify Source DB to migrate data from the specified DB in the source backup file. For example, if you enter 5, only data in DB5 will be migrated. To migrate all databases, do not specify Source DB.

- **Step 7** Enable **Multi-DB Proxy Cluster** if the source Redis is a multi-DB (**multi-db** set to **yes**) Proxy Cluster DCS Redis instance.
- **Step 8** Click **Add Backup** and select the backup files to be migrated.
- Step 9 In the Target Redis area, select the Target Redis Instance prepared in Prerequisites.
- Step 10 If the target Redis instance has a password, enter the password and click Test Connection to check whether the password is correct. If the instance is not password-protected, click Test Connection directly.
- Step 11 For Target DB, you can specify a DB in the target Redis to migrate data to. For example, if you enter 5, data will be migrated to DB5 of the target Redis. If you do not specify a DB, data will be migrated to a DB corresponding to the source DB.

D NOTE

- If the source Redis is multi-DB and the target is single-DB (DB0), either ensure that all source data is in DB0, or specify a source DB and set the target DB to **0**. Otherwise, migration will fail.
- For details about DB in DCS for Redis, see **Does DCS for Redis Support Multi-DB?**

#### Step 12 Click Next.

**Step 13** Confirm the migration task details and click **Submit**.

Go back to the data migration task list. After the migration is successful, the task status changes to **Successful**.

----End

# 11.5.3 Online Migration from Another Cloud Using Rump

Redis instances provided by some cloud service vendors do not allow **SLAVEOF**, **BGSAVE**, and **PSYNC** commands to be issued from Redis clients. As a result, rediscli, redis-shake, and other tools cannot be used to export data. Using the **KEYS** command may block Redis. Cloud service vendors usually only support downloading backup files. This method is suitable only for offline migration, featuring longer service interruption.

**Rump** is an open-source tool designed for migrating Redis data online. It supports migration between DBs of the same instance and between DBs of different instances. This section describes how to migrate another cloud to DCS by using Rump.

# Migration Principles

Rump uses the **SCAN** command to acquire keys and the **DUMP/RESTORE** command to get or set values.

Featuring time complexity O(1), **SCAN** is capable of quickly getting all keys. **DUMP/RESTORE** is used to read/write values independent from the key type.

Rump brings the following benefits:

- The SCAN command replaces the KEYS command to avoid blocking Redis.
- Any type of data can be migrated.

- **SCAN** and **DUMP/RESTORE** operations are pipelined, improving the network efficiency during data migration.
- No temporary file is involved, saving disk space.
- Buffered channels are used to optimize performance of the source server.

# **Notes and Constraints**

- The target cannot be a cluster DCS instance.
- To prevent migration command resolution errors, do not include special characters (#@:) in the instance password.
- Before data migration, suspend your services. If data is kept being written in during the migration, some keys might be lost.
- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

# Prerequisites

- If a target DCS Redis instance is not available, create one first. For details, see **Buying a DCS Redis Instance**.
- If you already have a DCS Redis instance, you do not need to create one again. For comparing migration data and reserving sufficient memory, you are advised to clear the instance data before the migration. For details, see Clearing DCS Instance Data. If any data exists on the target instance, duplicate data between the source and target is overwritten. If the data exists only on the target instance, the data will be retained.
- An Elastic Cloud Server (ECS) has been created. For details about how to create an ECS, see **Purchasing a Custom ECS**.

Select the same VPC, subnet, and security group as the DCS Redis Cluster instance, and bind EIPs to the ECS.

# Installing the Rump

- 1. Log in to the ECS.
- 2. Download Rump (release version).

On 64-bit Linux, run the following command:

wget https://github.com/stickermule/rump/releases/download/0.0.3/rump-0.0.3-linux-amd64;

 After decompression, run the following commands to add the execution permission: mv rump-0.0.3-linux-amd64 rump;

chmod +x rump;

# **Migrating Data**

Run the following command to migrate data:

rump -from {source\_redis\_address} -to {target\_redis\_address}

• {source\_redis\_address}

Source Redis instance address, in the format of redis:// [user:password@]host:port/db. **[user:password@]** is optional. If the instance is accessed in password-protected mode, you must specify the password in the RFC 3986 format. **user** can be omitted, but the colon (:) cannot be omitted. For example, the address may be **redis://:mypassword@192.168.0.45:6379/1**.

**db** is the sequence number of the database. If it is not specified, the default value is 0.

{target\_redis\_address}

Address of the target Redis instance, in the same format as the source.

In the following example, data in DB0 of the source Redis is migrated to the target Redis whose connection address is 192.168.0.153. \*\*\*\*\*\* stands for the password.

[root@ecs ~]# ./rump -from redis://127.0.0.1:6379/0 -to redis://:\*\*\*\*\*@192.168.0.153:6379/0 . Sync done. [root@ecs ~]#

# 11.5.4 Migrating from Another Cloud Online Using redisshake

redis-shake is an open-source Redis migration tool. Its **rump** mode allows you to obtain the full data of a source Redis using the **SCAN** command and write the data to a target Redis. This migration solution does not involve the **SYNC** or **PSYNC** command and can be widely used for migration between self-built Redis and cloud Redis.

This section describes how to use the **rump** mode of redis-shake to migrate the full Redis data of another cloud service vendor at a time online to Huawei Cloud DCS.





# Notes and Constraints

- The **rump** mode does not support incremental data migration. To keep data consistency, stop writing data to the source Redis before migration.
- This solution applies only to same-database mapping and does not apply to inter-database mapping.
- If the source Redis has multiple databases (there are databases other than DB0), and your Huawei Cloud DCS instance is Proxy Cluster, multi-DB must be enabled for the DCS instance. Otherwise, the migration will fail. (Single-DB Proxy Cluster instances do not support the **SELECT** command.)

- If the source Redis has multiple databases (there are databases other than DB0), and your Huawei Cloud DCS instance is Redis Cluster, this solution cannot be used. (Redis Cluster DCS instances support only DB0.)
- To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

# Prerequisites

- A DCS Redis instance has been created.
- An ECS has been created for running redis-shake. The ECS must use the same VPC as the Redis instance, and be bound to EIPs.

# Procedure

- **Step 1** Install Nginx on the Huawei Cloud ECS and the source forwarding server. The following describes how to install Nginx on an ECS running CentOS 7.x. The commands vary depending on the OS.
  - 1. Add Nginx to the Yum repository. sudo rpm -Uvh http://nginx.org/packages/centos/7/noarch/RPMS/nginx-releasecentos-7-0.el7.ngx.noarch.rpm
  - 2. Check whether Nginx has been added successfully. yum search nginx
  - 3. Install Nginx. sudo yum install -y nginx
  - 4. Install the stream module. yum install nginx-mod-stream --skip-broken
  - 5. Start Nginx and set it to run automatically upon system startup. sudo systemctl start nginx.service sudo systemctl enable nginx.service
  - 6. In the address box of a browser, enter the server address (the EIP of the ECS) to check whether Nginx is installed successfully.

If the following page is displayed, Nginx has been installed successfully.

|   | Welcome to <b>nginx</b> on Red Hat Enterprise Linux!   |
|---|--|
| This page is used to test the proper operation of the <b>nginx</b> HTTP server after it i   | has been installed. If you can read this page, it means that the web server installed at this site is working properly.  |
|   | Website Administrator  |
| This is the default iske keep bage that is distributed with <b>nginx</b> on Re<br>You should now put your content in a location of your choice and ed<br>For information on Red Hat Enterprise Linux, please visit the <u>Red Hat</u> | d Hat Enterprise Linux. It is located in /w//aww/aw/hast.<br>It the rest configuration directive in the <b>nginx</b> configuration file /sts/selss/selss.<br><u>Inc. website</u> . The documentation for Red Hat Enterprise Linux is <u>available on the Red Hat. Inc. website</u> . |
|   |  |

- Step 2 Add the source forwarding server to the whitelist of the source Redis.
- **Step 3** Configure a security group for the source forwarding server.
  - 1. Obtain the EIP of the Huawei Cloud ECS.
  - 2. In the inbound rule of the security group of the source forwarding server, add the EIP of the Huawei Cloud ECS, and open the port that Huawei Cloud ECS's requests come through. The following takes port 6379 as an example.
- **Step 4** Configure Nginx forwarding for the source forwarding server.
  - 1. Log in to the Linux source forwarding server and run the following commands to open and modify the configuration file:

```
cd /etc/nginx
vi nginx.conf
```

2. Example forwarding configuration:

```
stream {
	server {
	listen 6379;
	proxy_pass {source_instance_address}:{port};
	}
	}
```

**6379** is the listening port of the source forwarding server. *{source\_instance\_address}* and *{port}* are the connection address and port of the source Redis instance.

This configuration allows you to access the source Redis through the local listening port 6379 of the source forwarding server.

This configuration must be added exactly where it is shown in the following figure.

Figure 11-13 Configuration location



- 3. Restart Nginx. service nginx restart
- 4. Verify whether Nginx has been started. netstat -an|grep 6379

If the port is being listened, Nginx has been started successfully.

Figure 11-14 Verification result

tcp 0 0.0.0.0:6379 0.0.0.0:\* LISTEN

**Step 5** Configure Nginx forwarding for the Huawei Cloud ECS.

- Log in to the Linux ECS on Huawei Cloud and run the following commands to open and modify the configuration file: cd /etc/nginx vi nginx.conf
- 2. Configuration example:

```
stream {
    server {
        listen 6666;
        proxy_pass {source_ecs_address}:6379;
    }
}
```

**6666** is Huawei Cloud ECS's listening port, *{source\_ecs\_address}* is the public IP address of the source forwarding server, and **6379** is the listening port of the source forwarding server Nginx.

This configuration allows you to access the source forwarding server through the local listening port 6666 of the Huawei Cloud ECS. This configuration must be added exactly where it is shown in the following figure.

Figure 11-15 Configuration location

| <pre># Load dynamic modules. See /usr/share/doc/nginx/REA include /usr/share/nginx/modules/*.conf;</pre> | DME.dynamic. |
|--|--------------|
| events {   |              |
| worker_connections 1024;   |              |
| }  |              |
| stream {   |              |
| server {   |              |

- Restart Nginx. 3. service nginx restart
- 4. Verify whether Nginx has been started. netstat -an|grep 6666

If the port is being listened, Nginx has been started successfully.

| Figure 11-1 | 6 Verification | result |
|-------------|----------------|--------|
|-------------|----------------|--------|

0

0 0.0.0.0:6666 0.0.0.0:\* tcp LISTEN

**Step 6** Run the following command on the Huawei Cloud ECS to test the network connection of port 6666:

redis-cli -h {target\_ecs\_address} -p 6666 -a {password}

{target\_ecs\_address} is the EIP of the Huawei Cloud ECS, 6666 is the listening port of the Huawei Cloud ECS, and *{password}* is the source Redis password. If there is no password, leave it blank.

Figure 11-17 Connection example

```
[root@migrationtoolserver conf.d]# redis-cli -h
ЭK
:6666> info server
 Server
edis version:5.0.13
redis git shal:01fcc85a
redis_git_dirty:1
edis build id:97db56f84cd0ec69
edis mode:standalone
os:Linux
arch bits:64
multiplexing api:epoll
atomicvar api:atomic-builtin
gcc version:0.0.0
process id:102557
run id:a98007001c00368d619f772aaba236d704f585f9
tcp port:6379
uptime in seconds:899
uptime in days:0
hz:10
configured hz:10
lru clock:15186745
executable:
config file:
io threads active:0
19.0.1 :6666> info
```

**Step 7** Prepare the migration tool redis-shake.

- 1. Log in to the Huawei Cloud ECS.
- 2. Download redis-shake on the Huawei Cloud ECS. Version 2.0.3 is used as an example. You can use **other redis-shake versions** as required. wget https://github.com/tair-opensource/RedisShake/releases/download/release-v2.0.3-20200724/ redis-shake-v2.0.3.tar.gz
- 3. Decompress the redis-shake file. tar -xvf redis-shake-v2.0.3.tar.gz

**Step 8** Configure the redis-shake configuration file.

- 1. Go to the directory generated after the decompression. cd redis-shake-v2.0.3
- 2. Modify the **redis-shake.conf** configuration file. vim redis-shake.conf

Modify the source Redis configuration.

source.type

Type of the source Redis instance. Use **standalone** for single-node, master/standby, and Proxy Cluster, and **cluster** for cluster instances.

source.address

EIP of the Huawei Cloud ECS and the mapped port of the source forwarding server (Huawei Cloud ECS's listening port 6666). Separate the EIP and port number with a colon (:).

source.password\_raw

Password of the source Redis instance. If no password is set, you do not need to set this parameter.

Modify the target DCS configuration.

target.type

Type of the DCS Redis instance. Use **standalone** for single-node, master/ standby, and Proxy Cluster, and **cluster** for cluster instances.

- target.address
   Colon (:) separated connection address and port of the DCS Redis instance.
- target.password\_raw

Password of the DCS Redis instance. If no password is set, you do not need to set this parameter.

- 3. Press **Esc** to exit the editing mode and enter **:wq!**. Press **Enter** to save the configuration and exit the editing interface.
- **Step 9** Run the following command to start redis-shake and migrate data in the **rump** (online in full) mode:

./redis-shake.linux -conf redis-shake.conf -type rump

#### Figure 11-18 Migration process



Figure 11-19 Migration result



- **Step 10** After the migration is complete, use redis-cli to connect to the source and target Redis instances to check whether the data is complete.
  - 1. Connect to the source and target Redis instances, respectively. For details, see Access Using redis-cli.
  - 2. Run the **info keyspace** command to check the values of **keys** and **expires**.
  - 3. Calculate the differences between the values of **keys** and **expires** of the source Redis and the target Redis. If the differences are the same, the data is complete and the migration is successful.

**Step 11** After the verification is complete, you are advised to clear the redis-shake configuration in time.

----End

# 11.5.5 Backup Import from Another Cloud Using redis-shake

redis-shake is an open-source tool for migrating data online or offline (by importing backup files) between Redis Clusters. If the source Redis Cluster is deployed in another cloud, and online migration is not supported, you can migrate data by importing backup files.

If the source Redis and the target Redis cannot be connected, or the source Redis is deployed on other clouds, you can migrate data by importing backup files.

The following describes how to use redis-shake for backup migration to a DCS Redis Cluster instance.

# Notes and Constraints

To migrate to an instance with SSL enabled, disable the SSL setting first. For details, see **Transmitting DCS Redis Data with Encryption Using SSL**.

# Prerequisites

- A DCS Redis instance has been created. Note that the memory of the DCS Redis Cluster instance cannot be smaller than that of the source cluster.
- An ECS has been created for running redis-shake. The ECS must use the same VPC, subnet, and security group as the Redis instance.

# Procedure

- Access the target Redis instance using redis-cli. Obtain the IP address and port of the master node of the target instance. redis-cli -h {target\_redis\_address} -p {target\_redis\_port} -a {target\_redis\_password} cluster nodes
  - *{target\_redis\_address}*: connection address of the target DCS Redis instance.
  - *{target\_redis\_port}*: port of the target DCS Redis instance.
  - *{target\_redis\_password}*: password for connecting to the target DCS Redis instance.

In the command output similar to the following, obtain the IP addresses and ports of all masters.



- 2. Install redis-shake on the prepared Huawei Cloud ECS.
  - a. Log in to the Huawei Cloud ECS.
  - b. Download redis-shake on the Huawei Cloud ECS. Version 2.0.3 is used as an example. You can use **other redis-shake versions** as required. wget https://github.com/tair-opensource/RedisShake/releases/download/releasev2.0.3-20200724/redis-shake-v2.0.3.tar.gz

c. Decompress the redis-shake file. tar -xvf redis-shake-v2.0.3.tar.gz

#### **NOTE**

If the source Redis is deployed in the data center intranet, install redis-shake on the intranet server. Export data and then upload the data to the cloud server as instructed by the following steps.

- 3. Export the RDB file from the source Redis console. If the RDB file cannot be exported, contact customer service of the source.
- 4. Import the RDB file.
  - a. Import the RDB file (or files) to the cloud server. The cloud server must be connected to the target DCS instance.
  - b. Edit the redis-shake configuration file **redis-shake.conf**. vim redis-shake.conf

Add the following information about all the masters of the target: target.type = cluster # If there is no password, skip the following parameter. target.password\_raw = {target\_redis\_password} # IP addresses and port numbers of all masters of the target instance, which are separated by semicolons (;). target.address = {master1\_ip}:{master1\_port};{master2\_ip}:{master2\_port}...{masterN\_ip}: {masterN\_port} # List the RDB files to be imported, separated by semicolons (;). rdb.input = {local\_dump.0};{local\_dump.1};{local\_dump.2};{local\_dump.3}

Press **Esc** to exit the editing mode and enter **:wq!**. Press **Enter** to save the configuration and exit the editing interface.

c. Run the following command to import the RDB file to the target instance: ./redis-shake -type restore -conf redis-shake.conf

If the following information is displayed in the execution log, the backup file is imported successfully:

Enabled http stats, set status (incr), and wait forever.

5. Verify the migration.

After data synchronization, access the target Redis Cluster DCS instance using redis-cli. Run the **info** command to query the number of keys in the **Keyspace** section to confirm that data has been fully imported.

If the data has not been fully imported, run the **flushall** or **flushdb** command to clear the cached data in the instance, and synchronize data again.

# **12** Testing Instance Performance

# 12.1 Testing Redis Performance Using memtier\_benchmark

memtier\_benchmark is a command-line tool developed by Redis Labs. It can generate traffic in various modes and supports Redis. This tool provides multiple options and reporting features that can be easily used through the CLI. For details, visit https://github.com/RedisLabs/memtier\_benchmark.

This section describes how to use memtier-benchmark to test the performance of a DCS Redis instance when running command **SET** or **GET** in a high-concurrency scenario.

# **Test Procedure**

- Step 1 Create a DCS Redis instance.
- **Step 2** Create three ECSs and configure the same AZ, VPC, subnet, and security group for the ECSs and the instance.

**NOTE** 

Only one ECS is required for testing on a single-node or master/standby instance.

- Step 3 Install memtier\_benchmark on each ECS. CentOS 8.0 is used as an example.
  - 1. Preparations
    - a. Install the tools required for compilation.

yum install -y autoconf yum install -y automake yum install -y make yum install -y gcc-c++ yum install -y git

b. Enable the PowerTools repository.

dnf config-manager --set-enabled PowerTools

- c. Install the required dependencies.
  - yum install -y pcre-devel
  - yum install -y zlib-devel
  - yum install -y libmemcached-devel
  - yum install -y openssl-devel
- 2. Install the libevent library.

# yum install -y libevent-devel

- 3. Download, compile, and install the memtier\_benchmark library.
  - a. Create a folder in the root directory where the memtier\_benchmark library will be stored.

mkdir /env

b. Download the memtier\_benchmark source code.

cd /env

# git clone https://github.com/RedisLabs/memtier\_benchmark.git

c. Go to the directory where the source code is located.

# cd memtier\_benchmark

d. Compile the source code and generate the executable file **memtier\_benchmark**.

autoreconf -ivf

./configure

make

e. Install the tool in the system.

make install

4. Run the following command to check whether the installation is successful: If a parameter description of memtier\_benchmark is returned, the installation is successful.

# memtier\_benchmark --help

**Step 4** Run the following test command on all ECSs:

memtier\_benchmark -s {IP} -n {nreqs} -c {connect\_number} -t 4 -d {datasize}

# **NOTE**

Run memtier\_benchmark --cluster-mode -s {*IP*} -n {*nreqs*} -c {*connect\_number*} -t 4 -d {*datasize*} for a Redis Cluster instance.

# Reference values: -c *{connect\_number}*: 200; -n *{nreqs}*: 10,000,000; -d *{datasize}*: 32

- -s indicates the domain name or IP address of the instance.
- -t indicates the number of threads used in the benchmark test.
- -c indicates the number of client connections.
- -d indicates the size of a single data record in bytes.
- -n indicates the number of test packets.
- **Step 5** Repeat **Step 4** with different client connections to obtain the maximum QPS (Query per Second, number of read and write operations per second).

**Step 6** The sum of operations per second of all the three ECSs indicates the performance of the instance specification.

To test on a Redis Cluster instance, launch two benchmark tools on each ECS.

----End

# **12.2 Testing Redis Performance Using redis-benchmark**

The Redis client includes redis-benchmark, a performance testing utility that simulates N clients concurrently sending M number of query requests.

This section describes how to use redis-benchmark to test the performance of a DCS Redis instance when running command **SET** or **GET** in a high-concurrency scenario.

# **Test Procedure**

- **Step 1** Create a DCS Redis instance.
- **Step 2** Create three ECSs and configure the same AZ, VPC, subnet, and security group for the ECSs and the instance.

**NOTE** 

Only one ECS is required for testing on a single-node or master/standby instance.

- **Step 3** Install redis-benchmark on each ECS. The Redis server can be installed in either of the following ways and benchmark will be installed, too.
  - Method 1:
    - Download a Redis client. This example uses redis-6.0.9.
       wget http://download.redis.io/releases/redis-6.0.9.tar.gz
    - b. Decompress the client installation package.
       tar xzf redis-6.0.9.tar.gz
    - c. Go to the **src** directory of redis-6.0.9. **cd redis-6.0.9/src**
    - d. Compile the source code.
      - make

After the compilation is complete, the tool is stored in the **src** directory of **redis**-*x*.*x*.*x*.

- e. Check whether the redis-benchmark executable file exists.
  - ls

| [root       | src]‡            | ls                |                 |                   |                |              |
|-------------|------------------|-------------------|-----------------|-------------------|----------------|--------------|
| adlist.c    | conf ig.h        | geohash_helper.h  | lzfP.h          | rax.o             | scripting.o    | t_hash.c     |
| adlist.h    | config.o         | geohash_helper.o  | Makefile        | rdb.c             | sdsalloc.h     | t_hash.o     |
| adlist.o    | crc16.c          | geohash.o         | memtest.c       | rdb.h             | sds.c          | t_list.c     |
| ae.c        | crc16.0          | geo.o             | memtest.o       | rdb.o             | sds.h          | t_list.o     |
| ae_epoll.c  | crc64.c          | help.h            | mkreleasehdr.sh | redisassert.h     | sds.o          | t_set.c      |
| ae_evport.c | crc64.h          | hyperloglog.c     | module.c        | redis-benchmark   | sentinel.c     | t_set.o      |
| ae.h        | crc64.o          | hyperloglog.o     | module.o        | redis-benchmark.c | sentinel.o     | t_stream.c   |
| ae_kqueue.c | db.c             | intset.c          | modules         | redis-benchmark.o | server.c       | t_stream.o   |
| ae.o        | db.o             | intset.h          | multi.c         | redis-check-aof   | server.h       | t_string.c   |
| ae_select.c | debug.c          | intset.o          | multi.o         | redis-check-aof.c | server.o       | t_string.o   |
| anet.c      | debugmacro.h     | latency.c         | networking.c    | redis-check-aof.o | setproctitle.c | t_zset.c     |
| anet.h      | debug.o          | latency.h         | networking.o    | redis-check-rdb   | setproctitle.o | t_zset.o     |
| anet.o      | defrag.c         | latency.o         | notify.c        | redis-check-rdb.c | sha1.c         | util.c       |
| aof.c       | defrag.o         | lazyfree.c        | notify.o        | redis-check-rdb.o | sha1.h         | util.h       |
| aof .o      | dict.c           | lazyfree.o        | object.c        | redis-cli         | sha1.o         | util.o       |
| asciilogo.h | dict.h           | listpack.c        | object.o        | redis-cli.c       | siphash.c      | valgrind.sup |
| atomicvar.h | dict.o           | listpack.h        | pqsort.c        | redis-cli.o       | siphash.o      | version.h    |
| bio.c       | endianconv.c     | listpack_malloc.h | pqsort.h        | redismodule.h     | slowlog.c      | ziplist.c    |
| bio.h       | endianconv.h     | listpack.o        | pqsort.o        | redis-sentinel    | slowlog.h      | ziplist.h    |
| bio.o       | endianconv.o     | localtime.c       | pubsub.c        | redis-server      | slowlog.o      | ziplist.o    |
| bitops.c    | evict.c          | localtime.o       | pubsub.o        | redis-trib.rb     | solarisfixes.h | zipmap.c     |
| bitops.o    | evict.o          | lolwut5.c         | quicklist.c     | release.c         | sort.c         | zipmap.h     |
| blocked.c   | expire.c         | lolwut5.o         | quicklist.h     | release.h         | sort.o         | zipmap.o     |
| blocked.o   | expire.o         | lolwut.c          | quicklist.o     | release.o         | sparkline.c    | zmalloc.c    |
| childinfo.c | fmacros.h        | lolwut.o          | rand.c          | replication.c     | sparkline.h    | zmalloc.h    |
| childinfo.o | geo.c            | lzf_c.c           | rand.h          | replication.o     | sparkline.o    | zmalloc.o    |
| cluster.c   | geo.h            | lzf_c.o           | rand.o          | rio.c             | stream.h       |              |
| cluster.h   | geohash.c        | lzf_d.c           | rax.c           | rio.h             | syncio.c       |              |
| cluster.o   | geohash.h        | lzf_d.o           | rax.h           | rio.o             | syncio.o       |              |
| conf ig.c   | geohash_helper.c | lzf.h             | rax_malloc.h    | scripting.c       | testhelp.h     |              |

f. Install the tool in the system.

make install

• Method 2:

Install the Redis server matching the ECS OS. The following examples use Ubuntu and CentOS.

- Ubuntu
   sudo apt update
   sudo apt install redis-server
- CentOS
   sudo yum install epel-release
   sudo yum update
   sudo yum -y install redis

#### **Step 4** Run the following test command on all ECSs:

redis-benchmark -h *{IP}* -p *{Port}* -a *{password}* -n *{nreqs}* -r *{randomkeys}* -c *{connect\_number}* -d *{datasize}* -t *{command}* 

Reference values: -c *{connect\_number}*: 200; -n *{nreqs}*: 10,000,000; -r *{randomkeys}*: 1,000,000; -d *{datasize}*: 32

- -h: instance domain name or IP address
- -p: port of the instance. The default value is 6379.
- -a: password for connecting to the instance. This parameter is not required for password-free instances.
- -t Set of commands to be executed For example, to test only the set command, use -t set. To test the ping, get, and set commands, use -t ping,set,get. Use commas (,) to separate commands.
- -c number of client connections
- -d size of a single data record in bytes
- -n number of test packets
- -r number of random keys
- **Step 5** Repeat **Step 4** with different client connections to obtain the maximum QPS (Query per Second, number of read and write operations per second).
- **Step 6** The sum of operations per second of all the three ECSs indicates the performance of the instance specification.

To test on a Redis Cluster instance, launch two benchmark tools on each ECS.

### **NOTE**

- Add the --cluster parameter only when testing Redis Cluster instances using redisbenchmark.
- In a test for the maximum number of connections of a Redis Cluster instance, if the performance of the ECSs is insufficient, the program will exit or the error message "Cannot assign requested address" will be displayed when the number of connections reaches 10,000. In this case, check how many ECSs are used in the test. Prepare three ECSs and start three redis-benchmark processes on each ECS.

```
----End
```

# **Common redis-cli Options**

- -h <hostname>: host name of the server, which can be an IP address or a domain name.
- -p *<port>*: port of the server. The default value is **6379**.
- -a *<password>*: password for connecting to the server. This parameter is not required for password-free instances.
- -r <*repeat>*: number of times that a command is run.
- -n *<db>*: DB number. The default value is **0**.
- -c: cluster mode (with -ASK and -MOVED redirections).
- --latency: a loop where latency is measured continuously.
- --scan: scans the key space without blocking the Redis server. (By contrast, scanning using KEYS \* blocks Redis server).
- --eval <file>: sends the EVAL command using a Lua script.
- -x: reads the last parameter in stdin.
- --bigscan: scans big keys in the data set.
- --raw: forces raw data output from the hexadecimal format, such as \xe4\xb8.

# **Examples of Common redis-cli Commands**

- Connect to an instance:
   ./redis-cli -h {IP} -p 6379
- Connect to a specified DB:
  - ./redis-cli -h *{IP}* -p 6379 -n 10
- Connect to a Redis Cluster instance:
   ./redis-cli -h {*IP*} -p 6379 -c
- Test the latency (by sending the ping command):
   ./redis-cli -h {*IP*} -p 6379 --latency
- Scan for keys that match the specified pattern:
   ./redis-cli -h {/P} -p 6379 --scan --pattern '\*:12345\*'

# Common Options in redis-benchmark (redis-6.0.9)

 -h <hostname>: host name of the server, which can be an IP address or a domain name.

- -p *<port>*: port of the server. The default value is **6379**.
- -a <password>: password for connecting to the server. This parameter is not required for password-free instances.
- -c *<clients>*: number of concurrent connections. The default value is **50**.
- -n *<requests>*: total number of requests. The default value is **100000**.
- -d *<size*>: data size of the SET/GET value, in bytes. The default value is 2.
- --dbnum <db >: database number. The default value is 0.
- --threads <num>: multi-thread mode, which is supported only by redisbenchmark compiled in Redis 6.0. In pressure tests, the multi-thread mode outperforms the single-thread mode.
- --cluster: cluster mode (required only by Redis Cluster).
- -k <boolean>: 1=keep alive; 0=reconnect. The default value is 1, indicating that both pconnect and connect can be tested.
- -r <keyspacelen>: uses random keys for SET, GET, and INCR, and random values for SADD. keyspacelen indicates the number of keys to be added.
- -e: displays server errors to stdout.
- -q: displays only the number of queries per second.
- -l: runs tests in loops.
- -t <tests>: tests specified commands.
- -I : idle mode. Open *N* idle connections and wait.
- -P <*numreq*>: concurrent pipeline requests. The default value is 1.

For more information about redis-benchmark, visit https://redis.io/docs/latest/ operate/oss\_and\_stack/management/optimization/benchmarks/.

# **Examples of Common redis-benchmark Commands**

Test single-node, master/standby, read/write splitting, and Proxy Cluster instances:

./redis-benchmark -h *{IP address or domain name}* -p 6379 -a *{password}*--threads *{num}* -n *{ nreqs }* -r *{ randomkeys }* -c *{clients}* -d *{datasize}* -t *{command}* 

• Test Redis Cluster instances:

./redis-benchmark -h {IP address or domain name} -p 6379 -a
{password}--threads {num} -n { nreqs } -r { randomkeys } -c {clients} -d
{datasize} --cluster -t {command}

• Test connect:

./redis-benchmark -h *{IP address or domain name}* -p 6379 -a *{password}*--threads *{num}* -n *{ nreqs }* -r *{ randomkeys }* -c *{clients}* -d *{datasize}* -k 0 -t *{command}* 

• Test idle connections:

./redis-benchmark -h *{IP address or domain name}* -p 6379 -a *{pwd}* -c *{clients}* -I

# 12.3 Comparing redis-benchmark and memtier\_benchmark

| Tool                  | Memcached        | Setting<br>Read/Write<br>Ratio | Random<br>Payload | Setting<br>Timeout |
|-----------------------|------------------|--------------------------------|-------------------|--------------------|
| memtier_benc<br>hmark | Supported        | Supported                      | Supported         | Supported          |
| redis-<br>benchmark   | Not<br>supported | Not<br>supported               | Not supported     | Not<br>supported   |

# **12.4 Reference for a Redis Performance Test**

# 12.4.1 Test Data of Master/Standby DCS Redis 3.0 Instances

# **Test Environment**

- Redis instance specifications
  - Redis 3.0 | 8 GB | master/standby
  - Redis 3.0 | 32 GB | master/standby
- ECS flavors
   General computing-enhanced | c6.xlarge.2 | 4 vCPUs | 8 GB

# **Test Command**

redis-benchmark -h {IP} -p {Port} -a {password} -n {nreqs} -r {randomkeys} -c {connection} -d {datasize} -t {command}

Reference values: -c {connect\_number}: 1000; -n {nreqs}: 10,000,000; -r {randomkeys}: 1,000,000; -d {datasize}: 32

# **Test Result**

Table 12-1 Test result of running the SET command

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | 99.99 <sup>th</sup> -<br>Percentile<br>Latency (ms) | First 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|---|--|---|
| 8 GB                    | x86         | 1000                              | 107,65<br>7.69 | 20  | 23   | 27  |

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | 99.99 <sup>th</sup> -<br>Percentile<br>Latency (ms) | First 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|---|--|---|
|                         |             | 10,000                            | 72,750.<br>55  | 362   | 366  | 371   |
| 32 GB                   | x86         | 1000                              | 121,08<br>8.83 | 9   | 12   | 12  |
|                         |             | 10,000                            | 79,235.<br>53  | 203   | 204  | 267   |

Table 12-2 Test result of running the GET command

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | 99.99 <sup>th</sup> -<br>Percentile<br>Latency (ms) | First 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|---|--|---|
| 8 GB                    | x86         | 1000                              | 119,35<br>0.25 | 6   | 24   | 27  |
|                         |             | 10,000                            | 77,574.<br>7   | 152   | 358  | 365   |
| 32 GB                   | x86         | 1000                              | 124,65<br>0.98 | 16  | 17   | 17  |
|                         |             | 10,000                            | 81,991.<br>41  | 195   | 196  | 199   |

# D NOTE

DCS for Redis 3.0 does not support the Arm CPU architecture, so only x86-based instance test results are provided.

# 12.4.2 Test Data of Proxy Cluster DCS Redis 3.0 Instances

# **Test Environment**

- Redis instance specifications
  - Redis 3.0 | 64 GB | Proxy Cluster
- ECS flavors General computing-plus | c6.xlarge.2 | 4 vCPUs | 8 GB

# Test Command

redis-benchmark -h {IP} -p {Port} -a {password} -n {nreqs} -r {randomkeys} -c {connection} -d {datasize} -t {command}

Reference values: -c {connect\_number}: 1000; -n {nreqs}: 10,000,000; -r {randomkeys}: 1,000,000; -d {datasize}: 32

# **Test Result**

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | 99.99 <sup>th</sup> -<br>Percentile<br>Latency (ms) | First 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|---|--|---|
| 64 GB                   | x86         | 1000                              | 534,96<br>0.92 | 24  | 34   | 209   |
|                         |             | 10,000                            | 511,36<br>2.67 | 108   | 171  | 315   |

 Table 12-3 Test result of running the SET command

Table 12-4 Test result of running the GET command

| Redis<br>Cach<br>e Size | СРU<br>Туре | Concur<br>rent<br>Connec<br>tions | QPS            | 99.99 <sup>th</sup> -<br>Percentile<br>Latency (ms) | First 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|---|--|---|
| 64 GB                   | x86         | 1000                              | 584,66<br>9.15 | 23  | 31   | 82  |
|                         |             | 10,000                            | 533,17<br>8.04 | 144   | 184  | 370   |

# 

DCS for Redis 3.0 does not support the Arm CPU architecture, so only x86-based instance test results are provided.

# 12.4.3 Test Data of Master/Standby DCS Redis 4.0 or 5.0 Instances

# **Test Environment**

- Redis instance specifications
   Redis 4.0 or 5.0 | 8 GB | master/standby
   Redis 4.0 or 5.0 | 32 GB | master/standby
- ECS flavors General computing-enhanced | c6.2xlarge.2 | 8 vCPUs | 16 GB
- ECS image Ubuntu 18.04 server 64-bit

#### Test tool

A single ECS is used for the test. The test tool is redis-benchmark.

# Test Command

redis-benchmark -h {*IP*} -p {*Port*} -a {*password*} -n {*nreqs*} -r {*randomkeys*} -c {*connection*} -d {*datasize*} -t {*command*}

Reference values: -c *{connect\_number}*: 500; -n *{nreqs}*: 10,000,000; -r *{randomkeys}*: 1,000,000; -d *{datasize}*: 32; -t *{command}*: set

# Test Result

### **NOTE**

- The following test results are for reference only. The performance may vary depending on the site environment and network fluctuation.
- Certain cache sizes are taken for example in the following test results. For more information, see **DCS Instance Specifications**.
- QPS: Query per second, indicates number of read and write operations per second. Unit: count/second.
- Average Latency: Average latency of operations, in milliseconds.
- x<sup>th</sup> Percentile Latency: latency of x% of operations, in milliseconds. For example, if the value is 10 ms, 99.99<sup>th</sup> percentile latency indicates that 99.99% queries can be processed within 10 ms.

| Redis<br>Cach<br>e<br>Size | СРU<br>Туре | Concu<br>rrent<br>Conne<br>ctions | QPS            | 99.99 <sup>th</sup> -<br>Percenti<br>le<br>Latency<br>(ms) | First<br>100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last<br>100 <sup>th</sup> -<br>Percentil<br>e<br>Latency<br>(ms) | Average<br>Latency<br>(ms) |
|----------------------------|-------------|-----------------------------------|----------------|--|---|--|----------------------------|
| 8 GB                       | x86         | 500                               | 132,0<br>68.98 | 11   | 18  | 205  | 3.298                      |
|                            |             | 10,000                            | 82,38<br>6.58  | 171  | 178   | 263  | 69.275                     |
| 8 GB Arm                   | Arm         | 500                               | 94,81<br>1.89  | 10   | 12  | 13   | 3.476                      |
|                            |             | 10,000                            | 61,26<br>4.37  | 340  | 350   | 351  | 83.848                     |
| 32<br>GB                   | x86         | 500                               | 131,3<br>85.33 | 9.5  | 16  | 17   | 3.333                      |
|                            |             | 10,000                            | 82,27<br>5.41  | 157  | 162.18  | 162.43   | 62.105                     |
| 32<br>GB                   | Arm         | 500                               | 117,5<br>53.02 | 8  | 21  | 22   | 3.875                      |

#### Table 12-5 Test result of running the SET command

| Redis<br>Cach<br>e<br>Size | CPU<br>Type | Concu<br>rrent<br>Conne<br>ctions | QPS          | 99.99 <sup>th</sup> -<br>Percenti<br>le<br>Latency<br>(ms) | First<br>100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last<br>100 <sup>th</sup> -<br>Percentil<br>e<br>Latency<br>(ms) | Average<br>Latency<br>(ms) |
|----------------------------|-------------|-----------------------------------|--------------|--|---|--|----------------------------|
|                            |             | 10,000                            | 76,00<br>1.7 | 175  | 386   | 387  | 99.362                     |

Table 12-6 Test result of running the GET command

| Redis<br>Cach<br>e<br>Size | CPU<br>Type | Concu<br>rrent<br>Conne<br>ctions | QPS            | 99.99 <sup>th</sup> -<br>Percenti<br>le<br>Latency<br>(ms) | First<br>100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last<br>100 <sup>th</sup> -<br>Percentil<br>e Latency<br>(ms) | Average<br>Latency<br>(ms) |
|----------------------------|-------------|-----------------------------------|----------------|--|---|---|----------------------------|
| 8 GB                       | x86         | 500                               | 138,6<br>52.02 | 7  | 11  | 12  | 2.117                      |
|                            |             | 10,000                            | 82,71<br>0.94  | 123.7  | 281.6   | 282.9   | 61.078                     |
| 8 GB                       | Arm         | 500                               | 95,43<br>2.59  | 8.8  | 10  | 214   | 3.186                      |
|                            |             | 10,000                            | 60,98<br>4.16  | 217  | 337.15  | 337.92  | 83.321                     |
| 32<br>GB                   | x86         | 500                               | 139,1<br>13.02 | 6.6  | 10  | 11  | 2.119                      |
|                            |             | 10,000                            | 82,48<br>9.36  | 100  | 105.66  | 106   | 60.968                     |
| 32<br>GB                   | Arm         | 500                               | 139,0<br>41.45 | 6  | 10  | 11  | 2.487                      |
|                            |             | 10,000                            | 81,56<br>3.41  | 141  | 149   | 150   | 63                         |

# 12.4.4 Test Data of Proxy Cluster DCS Redis 4.0 or 5.0 Instances

# **Test Environment**

- Redis instance specifications
  - Redis 4.0 or 5.0 | 64 GB | 8 shards | Proxy Cluster

ECS flavors

General computing-enhanced | c6.xlarge.2 | 4 vCPUs | 8 GB

• Test tool

Three ECSs are used for concurrent tests. The test tool is memtier\_benchmark.

# **Test Command**

memtier\_benchmark --ratio= (1:0 and 0:1) -s {IP} -n {nreqs} -c {connect\_number} -t 4 -d {datasize}

Reference values: -c *{connect\_number}*: 1000; -n *{nreqs}*: 10,000,000; -d *{datasize}*: 32

# **Test Result**

### 

- The following test results are for reference only. The performance may vary depending on the site environment and network fluctuation.
- Certain cache sizes are taken for example in the following test results. For more information, see **DCS Instance Specifications**.
- QPS: Query per second, indicates number of read and write operations per second. Unit: count/second.
- Average Latency: Average latency of operations, in milliseconds.
- x<sup>th</sup> Percentile Latency: latency of x% of operations, in milliseconds. For example, if the value is 10 ms, 99.99<sup>th</sup> percentile latency indicates that 99.99% queries can be processed within 10 ms.

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS              | 95 <sup>th</sup> -<br>Percentile<br>Latency (ms) | 99.99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Maximum<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|------------------|--|--|----------------------------|
| 64 GB                   | x86         | 3000                              | 1,323,9<br>35.00 | 3.3  | 9.4  | 220                        |
|                         |             | 5000                              | 1,373,7<br>56.00 | 5.3  | 13   | 240                        |
|                         |             | 10,000                            | 1,332,0<br>74.00 | 11   | 26   | 230                        |
|                         |             | 80,000                            | 946,03<br>2.00   | 110  | 460  | 6800                       |
| 64 GB                   | Arm         | 3000                              | 837,86<br>4.92   | 5.8  | 16   | 78                         |
|                         |             | 5000                              | 763,60<br>9.69   | 10   | 29   | 240                        |
|                         |             | 10,000                            | 703,80<br>8.39   | 20   | 47   | 250                        |

#### Table 12-7 Test result of running the SET command
| Redis<br>Cach<br>e Size | СРU<br>Туре | Concur<br>rent<br>Connec<br>tions | QPS            | 95 <sup>th</sup> -<br>Percentile<br>Latency (ms) | 99.99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Maximum<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|--|--|----------------------------|
|                         |             | 80,000                            | 625,84<br>1.69 | 170  | 410  | 940                        |

#### Table 12-8 Test result of running the GET command

| Redis<br>Cach<br>e Size | СРU<br>Туре | Concur<br>rent<br>Connec<br>tions | QPS              | 95 <sup>th</sup> -<br>Percentile<br>Latency (ms) | 99.99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Maximum<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|------------------|--|--|----------------------------|
| 64 GB                   | x86         | 3000                              | 1,366,1<br>53.00 | 3.3  | 9.3  | 230                        |
|                         |             | 5000                              | 1,458,4<br>51.00 | 5.1  | 13   | 220                        |
|                         |             | 10,000                            | 1,376,3<br>99.00 | 11   | 29   | 440                        |
|                         |             | 80,000                            | 953,83<br>7.00   | 120  | 1300   | 2200                       |
| 64 GB                   | Arm         | 3000                              | 764,11<br>4.55   | 6.1  | 17   | 100                        |
|                         |             | 5000                              | 765,18<br>7.74   | 10   | 27   | 230                        |
|                         |             | 10,000                            | 731,31<br>0.95   | 20   | 47   | 250                        |
|                         |             | 80,000                            | 631,37<br>3.33   | 170  | 1300   | 1900                       |

# 12.4.5 Test Data of Redis Cluster DCS Redis 4.0 or 5.0 Instances

# **Test Environment**

- Redis instance specifications
  - Redis 4.0 or 5.0 | 32 GB | Redis Cluster
- ECS flavors General computing-enhanced | c6.xlarge.2 | 4 vCPUs | 8 GB
- Test tool

Three ECSs are used for concurrent tests. The test tool is memtier\_benchmark.

## **Test Command**

memtier\_benchmark --cluster-mode --ratio=(1:0 and 0:1) -s {IP} -n {nreqs} -c {connect\_number} -t 4 -d {datasize}

Reference values: -c *{connect\_number}*: 1000; -n *{nreqs}*: 10,000,000; -d *{datasize}*: 32

#### Test Result

#### 

- The following test results are for reference only. The performance may vary depending on the site environment and network fluctuation.
- Certain cache sizes are taken for example in the following test results. For more information, see DCS Instance Specifications.
- QPS: Query per second, indicates number of read and write operations per second. Unit: count/second.
- Average Latency: Average latency of operations, in milliseconds.
- x<sup>th</sup> Percentile Latency: latency of x% of operations, in milliseconds. For example, if the value is 10 ms, 99.99<sup>th</sup> percentile latency indicates that 99.99% queries can be processed within 10 ms.

| Redis<br>Cach<br>e Size | СРU<br>Туре | Concur<br>rent<br>Connec<br>tions | QPS            | 99.99 <sup>th</sup> -<br>Percentile<br>Latency (ms) | First 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|---|--|---|
| 32 GB                   | x86         | 1000                              | 371,78<br>0.2  | 5.6   | 6.3  | 44  |
|                         |             | 10,000                            | 256,07<br>3.11 | 90  | 220  | 460   |
| 32 GB                   | Arm         | 1000                              | 317,05<br>3.78 | 17  | 34   | 230   |
|                         |             | 10,000                            | 248,83<br>2.33 | 410   | 490  | 750   |

Table 12-9 Test result of running the SET command

Table 12-10 Test result of running the GET command

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | 99.99 <sup>th</sup> -<br>Percentile<br>Latency (ms) | First 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|---|--|---|
| 32 GB                   | x86         | 1000                              | 427,00<br>0.04 | 5.0   | 5.3  | 78  |

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | 99.99 <sup>th</sup> -<br>Percentile<br>Latency (ms) | First 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | Last 100 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|---|--|---|
|                         |             | 10,000                            | 302,15<br>9.03 | 63  | 220  | 460   |
| 32 GB                   | Arm         | 1000                              | 421,40<br>2.06 | 13  | 14   | 65  |
|                         |             | 10,000                            | 309,35<br>9.18 | 180   | 260  | 500   |

# 12.4.6 Test Data of Master/Standby DCS Redis 6.0 Instances

DCS Redis 6.0 basic edition instances support SSL. This section covers the performance tested with and without SSL enabled.

#### **Test Environment**

- Redis instance specifications
   Redis 6.0 | Basic edition | 8 GB | Master/Standby
   Redis 6.0 | Basic edition | 32 GB | Master/Standby
- ECS flavors General compute-plus | 8 vCPUs | 16 GiB | c7.2xlarge.2
- ECS image

Ubuntu 18.04 server 64-bit

Test tool

A single ECS is used for the test. The test tool is memtier\_benchmark.

## **Test Command**

SSL disabled:

./memtier\_benchmark -s {IP} -p {port} -c {connect\_number} -t {thread} -n allkeys --key-prefix="xxxx" --keyminimum=1 --key-maximum={max\_key} --key-pattern=P:P --ratio=1:0 -d {datasize}

Reference values: -c *{connect\_number}*: 1000, --key-maximum*{max\_key}*: 2000000, -d *{datasize}*: 32

#### SSL enabled:

./memtier\_benchmark -s {IP} -p {port} -c {connect\_number} -t {thread} -n allkeys --key-prefix="xxxx" --key-minimum=1 --key-maximum={max\_key} --key-pattern=P:P --ratio=1:0 -d {datasize} --tls --cacert ca.crt

Reference values: -c {connect\_number}: 1000, --key-maximum{max\_key}: 2000000, -d {datasize}: 32

# **Test Result**

#### D NOTE

- The following test results are for reference only. The performance may vary depending on the site environment and network fluctuation.
- Certain cache sizes are taken for example in the following test results. For more information, see **DCS Instance Specifications**.
- QPS: Query per second, indicates number of read and write operations per second. Unit: count/second.
- Average Latency: Average latency of operations, in milliseconds.
- x<sup>th</sup> Percentile Latency: latency of x% of operations, in milliseconds. For example, if the value is 10 ms, 99.99<sup>th</sup> percentile latency indicates that 99.99% queries can be processed within 10 ms.

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | Average<br>Latency (ms) | 99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | 99.9 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|-------------------------|---|---|
| 8 GB                    | 8 GB x86    | 500                               | 151,04<br>7.41 | 3.355                   | 6.175   | 12.223  |
|                         |             | 1000                              | 149,34<br>6.86 | 6.673                   | 11.711  | 31.743  |
| 32 GB                   | x86         | 500                               | 143,64<br>8.1  | 3.476                   | 5.215   | 13.055  |
|                         |             | 4000                              | 104,51<br>7.03 | 37.881                  | 139.263   | 175.103   |

Table 12-11 Test result of the SET command (SSL disabled)

 Table 12-12 Test result of the SET command (SSL enabled)

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS           | Average<br>Latency (ms) | 99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | 99.9 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|---------------|-------------------------|---|---|
| 8 GB                    | 8 GB x86    | 500                               | 86,827.<br>84 | 5.537                   | 8.575   | 9.535   |
|                         |             | 1000                              | 92,413.<br>99 | 10.055                  | 15.615  | 17.279  |
| 32 GB x86               | x86         | 500                               | 87,385.<br>5  | 5.584                   | 8.383   | 9.343   |
|                         |             | 4000                              | 50,813.<br>67 | 62.623                  | 100.863   | 104.959   |

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | Average<br>Latency (ms) | 99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | 99.9 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|-------------------------|---|---|
| 8 GB x86                | 500         | 180,41<br>3.66                    | 2.764          | 4.287                   | 11.583  |   |
|                         |             | 1000                              | 179,11<br>3.5  | 5.586                   | 8.959   | 29.823  |
| 32 GB x86               | x86         | 500                               | 175,26<br>8.86 | 2.848                   | 4.079   | 11.839  |
|                         |             | 4000                              | 134,75<br>5.17 | 29.161                  | 126.463   | 166.911   |

 Table 12-13 Test result of the GET command (SSL disabled)

**Table 12-14** Test result of the GET command (SSL enabled)

| Redis<br>Cach<br>e Size | СРU<br>Туре | Concur<br>rent<br>Connec<br>tions | QPS            | Average<br>Latency (ms) | 99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | 99.9 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|-------------------------|---|---|
| 8 GB                    | 8 GB x86    | 500                               | 113,63<br>7.22 | 4.316                   | 6.239   | 7.359   |
|                         |             | 1000                              | 105,50<br>4.55 | 8.962                   | 13.439  | 15.295  |
| 32 GB x86               | x86         | 500                               | 100,30<br>9.99 | 4.603                   | 6.559   | 6.943   |
|                         |             | 4000                              | 57,007.<br>69  | 55.052                  | 85.503  | 89.087  |

# 12.4.7 Test Data of Redis Cluster DCS Redis 6.0 Instances

DCS Redis 6.0 basic edition instances support SSL. This section covers the performance tested with and without SSL enabled.

## **Test Environment**

- Redis instance specifications Redis 6.0 | Basic edition | 32 GB | Redis Cluster
- ECS flavors General compute-plus | 8 vCPUs | 16 GiB | c7.2xlarge.2
- ECS image Ubuntu 18.04 server 64-bit

#### Test tool

Three ECSs are used for concurrent tests. The test tool is memtier\_benchmark.

#### **Test Command**

#### SSL disabled:

./memtier\_benchmark -s {IP} -p {port} -c {connect\_number} -t {thread} -n allkeys --key-prefix="xxxx" --keyminimum=1 --key-maximum={max\_key} --key-pattern=P:P --ratio=1:0 -d {datasize} --cluster-mode

Reference values: -c *{connect\_number}*: 1000, --key-maximum*{max\_key}*: 2000000, -d *{datasize}*: 32

#### SSL enabled:

./memtier\_benchmark -s {IP} -p {port} -c {connect\_number} -t {thread} -n allkeys --key-prefix="xxxx" --keyminimum=1 --key-maximum={max\_key} --key-pattern=P:P --ratio=1:0 -d {datasize} --cluster-mode --tls -cacert ca.crt

Reference values: -c *{connect\_number}*: 1000, --key-maximum*{max\_key}*: 2000000, -d *{datasize}*: 32

#### **Test Result**

#### D NOTE

- The following test results are for reference only. The performance may vary depending on the site environment and network fluctuation.
- Certain cache sizes are taken for example in the following test results. For more information, see **DCS Instance Specifications**.
- QPS: Query per second, indicates number of read and write operations per second. Unit: count/second.
- Average Latency: Average latency of operations, in milliseconds.
- x<sup>th</sup> Percentile Latency: latency of x% of operations, in milliseconds. For example, if the value is 10 ms, 99.99<sup>th</sup> percentile latency indicates that 99.99% queries can be processed within 10 ms.

| Redis<br>Cach<br>e Size | СРU<br>Туре | Concur<br>rent<br>Connec<br>tions | QPS            | Average<br>Latency (ms) | 99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | 99.9 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|-------------------------|---|---|
| 32 GB                   | x86         | 1000                              | 322,89<br>9.21 | 2.661                   | 4.319   | 8.511   |
|                         |             | 3000                              | 360,33<br>6.14 | 7.757                   | 13.055  | 29.439  |
|                         |             | 10,000                            | 330,37<br>8.22 | 29.411                  | 97.279  | 153,599   |

#### Table 12-15 Test result of the SET command (SSL disabled)

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | Average<br>Latency (ms) | 99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | 99.9 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|-------------------------|---|---|
| 32 GB                   | x86         | 1000                              | 238,30<br>7.26 | 3.603                   | 5.151   | 6.527   |
|                         |             | 3000                              | 185,45<br>5.62 | 13.196                  | 20.607  | 352.255   |
|                         |             | 10,000                            | 111,91<br>3.19 | 57.537                  | 96.767  | 121.343   |

 Table 12-16 Test result of the SET command (SSL enabled)

Table 12-17 Test result of the GET command (SSL disabled)

| Redis<br>Cach<br>e Size | СРU<br>Туре | Concur<br>rent<br>Connec<br>tions | QPS            | Average<br>Latency (ms) | 99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | 99.9 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|-------------------------|---|---|
| 32 GB                   | x86         | 1000                              | 450,42<br>2.66 | 1.875                   | 2.767   | 6.879   |
|                         |             | 3000                              | 432,45<br>0.2  | 6.451                   | 12.095  | 28.415  |
|                         |             | 10,000                            | 507,33<br>8.44 | 23.001                  | 95.231  | 176.127   |

 Table 12-18 Test result of the GET command (SSL enabled)

| Redis<br>Cach<br>e Size | CPU<br>Type | Concur<br>rent<br>Connec<br>tions | QPS            | Average<br>Latency (ms) | 99 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) | 99.9 <sup>th</sup> -<br>Percentile<br>Latency<br>(ms) |
|-------------------------|-------------|-----------------------------------|----------------|-------------------------|---|---|
| 32 GB                   | x86         | 1000                              | 274,06<br>6.16 | 3.076                   | 4.255   | 7.071   |
|                         |             | 3000                              | 201,06<br>3.51 | 11.743                  | 18.047  | 387.071   |
|                         |             | 10,000                            | 116,02<br>6.38 | 51.284                  | 84.479  | 136.191   |

# 12.4.8 Test Data of Redis Backup, Restoration, and Migration

#### **Test Environment**

- Redis instance specifications
  Redis 5.0 | 8 GB | master/standby
  Redis 5.0 | 32 GB | master/standby
  Redis 5.0 | 64 GB | Proxy Cluster (2 replicas | 8 shards | 8 GB per shard)
  Redis 5.0 | 256 GB | Proxy Cluster (2 replicas | 32 shards | 8 GB per shard)
  Redis 5.0 | 64 GB | Redis Cluster (2 replicas | 8 shards | 8 GB per shard)
  Redis 5.0 | 256 GB | Redis Cluster (2 replicas | 32 shards | 8 GB per shard)
- ECS flavors c6s.large.2 2 vCPUs | 4 GB

# **Test Command**

Run the following command on a 256 GB Proxy Cluster instance:

redis-benchmark - h {IP} -p{Port} -n 10000000 -r 10000000 -c 10000 -d 1024

Run the following command on a 256 GB Redis Cluster instance:

redis-benchmark - h {IP} -p{Port} -n 10000000 -r 10000000 -c 40000 -d 1024 -c

## **Test Result**

| Source<br>Instance<br>Type        | Source<br>Instance<br>Specifica<br>tions<br>(GB) | Target<br>Instance<br>Type        | Target<br>Instance<br>Specifica<br>tions<br>(GB) | Migratio<br>n Type        | Data<br>Volume<br>(GB) | Duration<br>(min) |
|-----------------------------------|--|-----------------------------------|--|---------------------------|------------------------|-------------------|
| Redis 5.0<br>  master/<br>standby | 8  | Redis 5.0<br>  master/<br>standby | 8  | Full +<br>incremen<br>tal | 7.78                   | 3                 |
| Redis 5.0<br>  master/<br>standby | 32   | Redis 5.0<br>  master/<br>standby | 32   | Full +<br>incremen<br>tal | 31.9                   | 17                |
| Redis 5.0<br>  Proxy<br>Cluster   | 64   | Redis 5.0<br>  Proxy<br>Cluster   | 64   | Full +<br>incremen<br>tal | 62.42                  | 7                 |
| Redis 5.0<br>  Redis<br>Cluster   | 64   | Redis 5.0<br>  Redis<br>Cluster   | 64   | Full +<br>incremen<br>tal | 57.69                  | 6                 |

Table 12-19 Migration

| Source<br>Instance<br>Type      | Source<br>Instance<br>Specifica<br>tions<br>(GB) | Target<br>Instance<br>Type      | Target<br>Instance<br>Specifica<br>tions<br>(GB) | Migratio<br>n Type        | Data<br>Volume<br>(GB) | Duration<br>(min) |
|---------------------------------|--|---------------------------------|--|---------------------------|------------------------|-------------------|
| Redis 5.0<br>  Proxy<br>Cluster | 256  | Redis 5.0<br>  Proxy<br>Cluster | 256  | Full +<br>incremen<br>tal | 241.48                 | 23                |
| Redis 5.0<br>  Redis<br>Cluster | 256  | Redis 5.0<br>  Redis<br>Cluster | 256  | Full +<br>incremen<br>tal | 240.21                 | 22                |

# Table 12-20 Backup

| Instance<br>Type                  | Instance<br>Specification<br>s (GB) | Backup Mode | Data Volume<br>(GB) | Duration<br>(min) |
|-----------------------------------|-------------------------------------|-------------|---------------------|-------------------|
| Redis 5.0  <br>master/<br>standby | 8                                   | RDB         | 7.78                | 2                 |
| Redis 5.0  <br>master/<br>standby | 32                                  | RDB         | 31.9                | 5                 |
| Redis 5.0  <br>Proxy Cluster      | 64                                  | RDB         | 62.42               | 9                 |
| Redis 5.0  <br>Proxy Cluster      | 256                                 | RDB         | 241.48              | 37                |
| Redis 5.0  <br>Redis Cluster      | 64                                  | RDB         | 57.69               | 9                 |
| Redis 5.0  <br>Redis Cluster      | 256                                 | RDB         | 255                 | 39                |
| Redis 5.0  <br>master/<br>standby | 8                                   | AOF         | 7.9                 | 2                 |
| Redis 5.0  <br>master/<br>standby | 32                                  | AOF         | 31.15               | 10                |
| Redis 5.0  <br>Proxy Cluster      | 64                                  | AOF         | 62.42               | 20                |
| Redis 5.0  <br>Proxy Cluster      | 256                                 | AOF         | 241.48              | 48                |

| Instance<br>Type             | Instance<br>Specification<br>s (GB) | Backup Mode | Data Volume<br>(GB) | Duration<br>(min) |
|------------------------------|-------------------------------------|-------------|---------------------|-------------------|
| Redis 5.0  <br>Redis Cluster | 64                                  | AOF         | 57.69               | 19                |
| Redis 5.0  <br>Redis Cluster | 256                                 | AOF         | 255                 | 51                |

#### Table 12-21 Restoration

| Instance<br>Type                  | Instance<br>Specification<br>s (GB) | Restoration<br>Mode | Data Volume<br>(GB) | Duration<br>(min) |
|-----------------------------------|-------------------------------------|---------------------|---------------------|-------------------|
| Redis 5.0  <br>master/<br>standby | 8                                   | RDB                 | 7.9                 | 2                 |
| Redis 5.0  <br>master/<br>standby | 32                                  | RDB                 | 31.15               | 6                 |
| Redis 5.0  <br>Proxy Cluster      | 64                                  | RDB                 | 62.42               | 10                |
| Redis 5.0  <br>Proxy Cluster      | 256                                 | RDB                 | 246                 | 42                |
| Redis 5.0  <br>Redis Cluster      | 64                                  | RDB                 | 57.69               | 10                |
| Redis 5.0  <br>Redis Cluster      | 256                                 | RDB                 | 255                 | 40                |
| Redis 5.0  <br>master/<br>standby | 8                                   | AOF                 | 7.9                 | 3                 |
| Redis 5.0  <br>master/<br>standby | 32                                  | AOF                 | 31.15               | 10                |
| Redis 5.0  <br>Proxy Cluster      | 64                                  | AOF                 | 62.42               | 10                |
| Redis 5.0  <br>Proxy Cluster      | 256                                 | AOF                 | 246                 | 46                |
| Redis 5.0  <br>Redis Cluster      | 64                                  | AOF                 | 57.69               | 10                |

| Instance<br>Type             | Instance<br>Specification<br>s (GB) | Restoration<br>Mode | Data Volume<br>(GB) | Duration<br>(min) |
|------------------------------|-------------------------------------|---------------------|---------------------|-------------------|
| Redis 5.0  <br>Redis Cluster | 256                                 | AOF                 | 255                 | 43                |

# **13** Applying for More DCS Quotas

## What Is Quota?

A quota is a limit on the quantity or capacity of a certain type of service resources that you can use, for example, the maximum number of DCS instances that you can create and the maximum amount of memory that you can use.

If a quota cannot meet your needs, apply for a higher quota.

#### How Do I View My Quota?

- 1. Log in to the management console.
- 2. Click  $\bigcirc$  in the upper left corner of the console and select the region where your instance is located.
- In the upper right corner of the page, choose Resources > My Quotas. The Service Quota page is displayed.

#### Figure 13-1 My Quotas



4. On the **Service Quota** page, view the used and total quotas of resources. If a quota cannot meet your needs, apply for a higher quota by performing the following operations.

## How Do I Increase My Quota?

1. Log in to the management console.

 In the upper right corner of the page, choose Resources > My Quotas. The Service Quota page is displayed.

Figure 13-2 My Quotas



- 3. Click Increase Quota.
- 4. On the Create Service Ticket page, set the parameters.

In the **Problem Description** area, enter the required quota and the reason for the quota adjustment.

5. Read the agreements and confirm that you agree to them, and then click **Submit**.

# **14** Viewing Monitoring Metrics and Configuring Alarms

Cloud Eye is a secure, scalable monitoring platform. It monitors DCS metrics, and sends notifications if alarms are triggered or events occur.

# **14.1 DCS Metrics**

# Introduction

This section describes DCS metrics reported to Cloud Eye as well as their namespaces and dimensions. You can use the Cloud Eye console or call **APIs** to query the DCS metrics and alarms.

Different types of instances are monitored on different dimensions.

| Instance<br>Type        | Instance<br>Monitoring  | Redis Server<br>Monitoring                                     | Proxy Monitoring                         |
|-------------------------|---|--|--|
| Single-<br>node         | Supported<br>The monitoring on<br>the instance<br>dimension is<br>conducted on the<br>Redis Server. | N/A  | N/A                                      |
| Master/<br>standby      | Supported<br>The master node is<br>monitored.   | Supported<br>The master and<br>standby nodes are<br>monitored. | N/A                                      |
| Read/write<br>splitting | Supported<br>The master node is<br>monitored.   | Supported<br>The master and<br>standby nodes are<br>monitored. | Supported<br>Each proxy is<br>monitored. |

 Table 14-1
 Monitoring dimensions for different instance types

| Instance<br>Type | Instance<br>Monitoring  | Redis Server<br>Monitoring               | Proxy Monitoring                         |
|------------------|---|--|--|
| Proxy<br>Cluster | Supported<br>The monitoring<br>data is the<br>aggregated master<br>node data. | Supported<br>Each shard is<br>monitored. | Supported<br>Each proxy is<br>monitored. |
| Redis<br>Cluster | Supported<br>The monitoring<br>data is the<br>aggregated master<br>node data. | Supported<br>Each shard is<br>monitored. | N/A                                      |

# Namespace

SYS.DCS

# **DCS Redis 3.0 Instance Metrics**

#### **NOTE**

- DCS for Redis 3.0 is no longer provided. You can use DCS for Redis 4.0 or later instead.
- **Dimensions** lists the metric dimensions.

| Metric ID                  | Metri<br>c<br>Nam<br>e                      | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t   | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|----------------------------|---|---|--------------------|------------|--------------------------------|--|---|
| cpu_usag<br>e              | CPU<br>Usag<br>e                            | The monitored<br>object's maximum<br>CPU usage among<br>multiple sampling<br>values in a<br>monitoring period<br>For a single-node<br>or master/standby<br>instance, this<br>metric indicates<br>the CPU usage of<br>the master node.<br>For a Proxy<br>Cluster instance,<br>this metric<br>indicates the<br>average value of<br>all proxies. | 0-<br>100%         | %          | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| memory_<br>usage           | Mem<br>ory<br>Usag<br>e                     | Memory<br>consumed by the<br>monitored object<br><b>NOTICE</b><br>The memory usage<br>does not include<br>the usage of<br>reserved memory.  | 0-<br>100%         | %          | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| net_in_thr<br>oughput      | Netw<br>ork<br>Input<br>Throu<br>ghput      | Inbound<br>throughput per<br>second on a port   | ≥ 0                | byte<br>/s | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| net_out_t<br>hroughpu<br>t | Netw<br>ork<br>Outp<br>ut<br>Throu<br>ghput | Outbound<br>throughput per<br>second on a port  | ≥ 0                | byte<br>/s | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

Table 14-2 DCS Redis 3.0 instance metrics

| Metric ID                       | Metri<br>c<br>Nam<br>e                      | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------------------|---|---|--------------------|----------|--------------------------------|--|---|
| connecte<br>d_clients           | Conn<br>ected<br>Client<br>s                | Number of<br>connected clients.<br>Includes<br>connections<br>established for<br>system<br>monitoring,<br>configuration<br>synchronization,<br>and services.<br>Excludes<br>connections from<br>replicas. | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| client_lon<br>gest_out_l<br>ist | Client<br>Long<br>est<br>Outp<br>ut<br>List | Longest output<br>list among current<br>client connections  | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| client_big<br>gest_in_b<br>uf   | Client<br>Bigge<br>st<br>Input<br>Buf       | Maximum input<br>data length<br>among current<br>client connections   | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| blocked_c<br>lients             | Block<br>ed<br>Client<br>s                  | Number of clients<br>suspended by<br>block operations<br>such as BLPOP,<br>BRPOP, and<br>BRPOPLPUSH   | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| used_me<br>mory                 | Used<br>Mem<br>ory                          | Number of bytes<br>used by the Redis<br>server  | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID                          | Metri<br>c<br>Nam<br>e                       | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|------------------------------------|--|--|--------------------|----------|--------------------------------|--|---|
| used_me<br>mory_rss                | Used<br>Mem<br>ory<br>RSS                    | Resident set size<br>(RSS) memory<br>that the Redis<br>server has used,<br>which is the<br>memory that<br>actually resides in<br>the memory,<br>including all stack<br>and heap memory<br>but not swapped-<br>out memory | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| used_me<br>mory_pea<br>k           | Used<br>Mem<br>ory<br>Peak                   | Peak memory<br>consumed by<br>Redis since the<br>Redis server last<br>started  | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| used_me<br>mory_lua                | Used<br>Mem<br>ory<br>Lua                    | Number of bytes<br>used by the Lua<br>engine   | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| memory_f<br>rag_ratio              | Mem<br>ory<br>Frag<br>ment<br>ation<br>Ratio | Current memory<br>fragmentation,<br>which is the ratio<br>between<br>used_memory_rs<br>s/used_memory.  | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| total_con<br>nections_r<br>eceived | New<br>Conn<br>ectio<br>ns                   | Number of<br>connections<br>received during<br>the monitoring<br>period  | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| total_com<br>mands_pr<br>ocessed   | Com<br>mand<br>s<br>Proce<br>ssed            | Number of<br>commands<br>processed during<br>the monitoring<br>period  | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID                         | Metri<br>c<br>Nam<br>e              | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t  | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------------------|-------------------------------------|--|--------------------|-----------|--------------------------------|--|---|
| instantan<br>eous_ops             | Ops<br>per<br>Secon<br>d            | Number of<br>commands<br>processed per<br>second   | ≥ 0                | N/A       | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| total_net_<br>input_byt<br>es     | Netw<br>ork<br>Input<br>Bytes       | Number of bytes<br>received during<br>the monitoring<br>period   | ≥ 0                | byte      | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| total_net_<br>output_by<br>tes    | Netw<br>ork<br>Outp<br>ut<br>Bytes  | Number of bytes<br>sent during the<br>monitoring period  | ≥ 0                | byte      | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| instantan<br>eous_inpu<br>t_kbps  | Input<br>Flow                       | Instantaneous<br>input traffic   | ≥ 0                | KiB/<br>s | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| instantan<br>eous_out<br>put_kbps | Outp<br>ut<br>Flow                  | Instantaneous<br>output traffic  | ≥ 0                | KiB/<br>s | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| rejected_c<br>onnection<br>s      | Reject<br>ed<br>Conn<br>ectio<br>ns | Number of<br>connections that<br>have exceeded<br>maxclients and<br>been rejected<br>during the<br>monitoring period | ≥ 0                | N/A       | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| expired_k<br>eys                  | Expire<br>d<br>Keys                 | Number of keys<br>that have expired<br>and been deleted<br>during the<br>monitoring period                           | ≥ 0                | N/A       | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| evicted_k<br>eys                  | Evicte<br>d<br>Keys                 | Number of keys<br>that have been<br>evicted and<br>deleted during the<br>monitoring period                           | ≥ 0                | N/A       | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID                 | Metri<br>c<br>Nam<br>e                    | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------------|---|---|--------------------|----------|--------------------------------|--|---|
| keyspace_<br>hits         | Keysp<br>ace<br>Hits                      | Number of<br>successful lookups<br>of keys in the<br>main dictionary<br>during the<br>monitoring period   | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| keyspace_<br>misses       | Keysp<br>ace<br>Misse<br>s                | Number of failed<br>lookups of keys in<br>the main<br>dictionary during<br>the monitoring<br>period   | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| pubsub_c<br>hannels       | PubS<br>ub<br>Chan<br>nels                | Number of<br>Pub/Sub channels   | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| pubsub_p<br>atterns       | PubS<br>ub<br>Patter<br>ns                | Number of<br>Pub/Sub patterns   | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| keyspace_<br>hits_perc    | Hit<br>Rate                               | Ratio of the<br>number of Redis<br>cache hits to the<br>number of<br>lookups.<br>Calculation:<br>keyspace_hits/<br>(keyspace_hits +<br>keyspace_misses) | 0–<br>100%         | %        | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| command<br>_max_del<br>ay | Maxi<br>mum<br>Com<br>mand<br>Laten<br>cy | Maximum latency<br>of commands  | ≥ 0                | ms       | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID             | Metri<br>c<br>Nam<br>e                 | Metric<br>Description   | Value<br>Rang<br>e                                     | Uni<br>t  | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)                           | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------|--|---|--|-----------|--------------------------------|--|---|
| auth_erro<br>rs       | Authe<br>nticat<br>ion<br>Failur<br>es | Number of failed<br>authentications   | ≥ 0  | Cou<br>nt | N/A                            | Single-<br>node or<br>master/<br>standby<br>DCS<br>Redis<br>instance | 1<br>minut<br>e                                       |
| is_slow_lo<br>g_exist | Slow<br>Quer<br>y<br>Logs              | Existence of slow<br>query logs in the<br>instance<br>NOTE<br>Slow queries<br>caused by the<br>MIGRATE,<br>SLAVEOF,<br>CONFIG, BGSAVE,<br>and<br>BGREWRITEAOF<br>commands are not<br>counted. | <ul> <li>1:<br/>ye<br/>s</li> <li>0:<br/>no</li> </ul> | N/A       | N/A                            | Single-<br>node or<br>master/<br>standby<br>DCS<br>Redis<br>instance | 1<br>minut<br>e                                       |
| keys                  | Keys                                   | Number of keys in<br>Redis  | ≥ 0  | N/A       | N/A                            | Single-<br>node or<br>master/<br>standby<br>DCS<br>Redis<br>instance | 1<br>minut<br>e                                       |

# **DCS Redis 4.0 and Later Instance Metrics**

#### D NOTE

- **Dimensions** lists the metric dimensions.
- The monitoring data is the aggregated master node data.
- Some metrics are aggregated from the master and replica nodes. For details, see "Metric Description" in Table 14-3.

| Metric ID                          | Metri<br>c<br>Nam<br>e                    | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|------------------------------------|---|--|--------------------|----------|--------------------------------|---|---|
| cpu_usag<br>e                      | CPU<br>Usag<br>e                          | The monitored<br>object's maximum<br>CPU usage among<br>multiple sampling<br>values in a<br>monitoring period  | 0-100              | %        | N/A                            | Single-<br>node,<br>master/<br>standby,<br>or read/<br>write<br>splitting<br>DCS<br>Redis<br>instance | 1<br>minut<br>e                                       |
| cpu_avg_<br>usage                  | Avera<br>ge<br>CPU<br>Usag<br>e           | The monitored<br>object's average<br>CPU usage of<br>multiple sampling<br>values in a<br>monitoring period   | 0-100              | %        | N/A                            | Single-<br>node,<br>master/<br>standby,<br>or read/<br>write<br>splitting<br>DCS<br>Redis<br>instance | 1<br>minut<br>e                                       |
| command<br>_max_del<br>ay          | Maxi<br>mum<br>Com<br>mand<br>Laten<br>cy | Maximum latency<br>of commands   | ≥ 0                | ms       | N/A                            | DCS<br>Redis<br>instance  | 1<br>minut<br>e                                       |
| total_con<br>nections_r<br>eceived | New<br>Conn<br>ectio<br>ns                | Number of<br>connections<br>received during<br>the monitoring<br>period. Includes<br>connections from<br>replicas and<br>established for<br>system<br>monitoring,<br>configuration<br>synchronization,<br>and services | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance  | 1<br>minut<br>e                                       |

 Table 14-3 DCS Redis 4.0 and later instance metrics

| Metric ID              | Metri<br>c<br>Nam<br>e               | Metric<br>Description  | Value<br>Rang<br>e                                     | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|------------------------|--------------------------------------|--|--|----------|--------------------------------|--|---|
| is_slow_lo<br>g_exist  | Slow<br>Quer<br>y<br>Logs            | Existence of slow<br>query logs in the<br>instance<br>NOTE<br>Slow queries<br>caused by the<br>MIGRATE,<br>SLAVEOF,<br>CONFIG, BGSAVE,<br>and<br>BGREWRITEAOF<br>commands are not<br>counted.  | <ul> <li>1:<br/>ye<br/>s</li> <li>0:<br/>no</li> </ul> | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| memory_<br>usage       | Mem<br>ory<br>Usag<br>e              | Memory<br>consumed by the<br>monitored object  | 0-100  | %        | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| expires                | Keys<br>With<br>an<br>Expir<br>ation | Number of keys<br>with an expiration<br>in Redis   | ≥ 0  | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| keyspace_<br>hits_perc | Hit<br>Rate                          | Ratio of the<br>number of Redis<br>cache hits to the<br>number of<br>lookups.<br>Calculation:<br>keyspace_hits/<br>(keyspace_hits +<br>keyspace_misses)<br>Aggregated from<br>the master and<br>replica nodes. If<br>no read command<br>is performed<br>within a<br>monitoring<br>period, the ratio is<br>0. | 0-100  | %        | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID                        | Metri<br>c<br>Nam<br>e                     | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|----------------------------------|--|--|--------------------|----------|--------------------------------|--|---|
| used_me<br>mory                  | Used<br>Mem<br>ory                         | Total number of<br>bytes used by the<br>Redis server   | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| used_me<br>mory_dat<br>aset      | Used<br>Mem<br>ory<br>Datas<br>et          | Dataset memory<br>that the Redis<br>server has used  | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| used_me<br>mory_dat<br>aset_perc | Used<br>Mem<br>ory<br>Datas<br>et<br>Ratio | Percentage of<br>dataset memory<br>that server has<br>used<br>Aggregated from<br>the master and<br>replica nodes.  | 0-100              | %        | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| used_me<br>mory_rss              | Used<br>Mem<br>ory<br>RSS                  | Resident set size<br>(RSS) memory<br>that the Redis<br>server has used,<br>which is the<br>memory that<br>actually resides in<br>the memory,<br>including all stack<br>and heap memory<br>but not swapped-<br>out memory | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| instantan<br>eous_ops            | Ops<br>per<br>Secon<br>d                   | Number of<br>commands<br>processed per<br>second   | ≥ 0                | N/A      | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID             | Metri<br>c<br>Nam<br>e       | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------|------------------------------|---|--------------------|-------------|--------------------------------|--|---|
| keyspace_<br>misses   | Keysp<br>ace<br>Misse<br>s   | Number of failed<br>lookups of keys in<br>the main<br>dictionary during<br>the monitoring<br>period<br>Aggregated from<br>the master and<br>replica nodes.  | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| keys                  | Keys                         | Number of keys in<br>Redis  | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| blocked_c<br>lients   | Block<br>ed<br>Client<br>s   | Number of clients<br>suspended by<br>block operations   | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| connecte<br>d_clients | Conn<br>ected<br>Client<br>s | Number of<br>connected clients.<br>Includes<br>connections<br>established for<br>system<br>monitoring,<br>configuration<br>synchronization,<br>and services.<br>Excludes<br>connections from<br>replicas. | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| del                   | DEL                          | Number of <b>DEL</b><br>commands<br>processed per<br>second   | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID        | Metri<br>c<br>Nam<br>e | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|------------------|------------------------|---|--------------------|-------------|--------------------------------|--|---|
| evicted_k<br>eys | Evicte<br>d<br>Keys    | Number of keys<br>that have been<br>evicted and<br>deleted during the<br>monitoring period<br>Aggregated from<br>the master and<br>replica nodes. | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| expire           | EXPIR<br>E             | Number of<br><b>EXPIRE</b><br>commands<br>processed per<br>second   | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| expired_k<br>eys | Expire<br>d<br>Keys    | Number of keys<br>that have expired<br>and been deleted<br>during the<br>monitoring period<br>Aggregated from<br>the master and<br>replica nodes. | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| get              | GET                    | Number of <b>GET</b><br>commands<br>processed per<br>second<br>Aggregated from<br>the master and<br>replica nodes.                                | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| hdel             | HDEL                   | Number of <b>HDEL</b><br>commands<br>processed per<br>second  | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| hget             | HGET                   | Number of <b>HGET</b><br>commands<br>processed per<br>second<br>Aggregated from<br>the master and<br>replica nodes.                               | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID                         | Metri<br>c<br>Nam<br>e                       | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------------------|--|---|--------------------|-------------|--------------------------------|--|---|
| hmget                             | HMG<br>ET                                    | Number of<br>HMGET<br>commands<br>processed per<br>second<br>Aggregated from<br>the master and<br>replica nodes.    | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| hmset                             | HMS<br>ET                                    | Number of<br>HMSET<br>commands<br>processed per<br>second   | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| hset                              | HSET   | Number of <b>HSET</b><br>commands<br>processed per<br>second  | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| instantan<br>eous_inpu<br>t_kbps  | Input<br>Flow                                | Instantaneous<br>input traffic  | ≥ 0                | KiB/<br>s   | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| instantan<br>eous_out<br>put_kbps | Outp<br>ut<br>Flow                           | Instantaneous<br>output traffic   | ≥ 0                | KiB/<br>s   | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| memory_f<br>rag_ratio             | Mem<br>ory<br>Frag<br>ment<br>ation<br>Ratio | Current memory<br>fragmentation   | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| mget                              | MGET   | Number of <b>MGET</b><br>commands<br>processed per<br>second<br>Aggregated from<br>the master and<br>replica nodes. | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID                | Metri<br>c<br>Nam<br>e     | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|--------------------------|----------------------------|---|--------------------|-------------|--------------------------------|--|---|
| mset                     | MSET                       | Number of <b>MSET</b><br>commands<br>processed per<br>second  | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| pubsub_c<br>hannels      | PubS<br>ub<br>Chan<br>nels | Number of<br>Pub/Sub channels   | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| pubsub_p<br>atterns      | PubS<br>ub<br>Patter<br>ns | Number of<br>Pub/Sub patterns   | ≥ 0                | N/A         | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| set                      | SET                        | Number of <b>SET</b><br>commands<br>processed per<br>second   | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| used_me<br>mory_lua      | Used<br>Mem<br>ory<br>Lua  | Number of bytes<br>used by the Lua<br>engine  | ≥ 0                | byte        | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| used_me<br>mory_pea<br>k | Used<br>Mem<br>ory<br>Peak | Peak memory<br>consumed by<br>Redis since the<br>Redis server last<br>started                                       | ≥ 0                | byte        | 102<br>4(IE<br>C)              | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| sadd                     | Sadd                       | Number of <b>SADD</b><br>commands<br>processed per<br>second  | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |
| smember<br>s             | Sme<br>mber<br>s           | Number of<br>SMEMBERS<br>commands<br>processed per<br>second<br>Aggregated from<br>the master and<br>replica nodes. | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                   | 1<br>minut<br>e                                       |

| Metric ID           | Metri<br>c<br>Nam<br>e       | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------|------------------------------|--|--------------------|-------------|--------------------------------|---|---|
| scan                | SCAN                         | Number of SCAN<br>operations per<br>second   | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                    | 1<br>minut<br>e                                       |
| setex               | SETE<br>X                    | Number of SETEX operations per second  | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | DCS<br>Redis<br>instance                    | 1<br>minut<br>e                                       |
| rx_control<br>led   | Flow<br>Contr<br>ol<br>Times | Number of times<br>that client<br>requests are<br>controlled in a<br>period. This<br>metric is<br>incremented by 1<br>each time a client<br>request is<br>controlled.<br>If the value is<br>greater than 0,<br>the consumed<br>bandwidth<br>exceeds the upper<br>limit and flow<br>control is<br>triggered on a<br>node. The node<br>suspends client<br>commands<br>temporarily. | ≥ 0                | Cou<br>nt   | N/A                            | DCS<br>Redis<br>instance                    | 1<br>minut<br>e                                       |
| bandwidt<br>h_usage | Band<br>width<br>Usag<br>e   | Percentage of the<br>used bandwidth<br>to the maximum<br>bandwidth limit   | 0-200              | %           | N/A                            | DCS<br>Redis<br>instance                    | 1<br>minut<br>e                                       |
| command<br>_max_rt  | Maxi<br>mum<br>Laten<br>cy   | Maximum delay<br>from when the<br>node receives<br>commands to<br>when it responds   | ≥ 0                | μs          | N/A                            | Single-<br>node<br>DCS<br>Redis<br>instance | 1<br>minut<br>e                                       |

| Metric ID          | Metri<br>c<br>Nam<br>e       | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)                           | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|--------------------|------------------------------|--|--------------------|----------|--------------------------------|--|---|
| command<br>_avg_rt | Avera<br>ge<br>Laten<br>cy   | Average delay<br>from when the<br>node receives<br>commands to<br>when it responds | ≥ 0                | μs       | N/A                            | Single-<br>node<br>DCS<br>Redis<br>instance                          | 1<br>minut<br>e                                       |
| used_stor<br>age   | Used<br>Stora<br>ge<br>Space | Storage space<br>that has been<br>used.  | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Redis<br>professio<br>nal<br>(storage)<br>edition<br>instance | 1<br>minut<br>e                                       |
| storage_u<br>sage  | Stora<br>ge<br>Usag<br>e     | Percentage of<br>used storage.   | 0-100              | %        | N/A                            | DCS<br>Redis<br>professio<br>nal<br>(storage)<br>edition<br>instance | 1<br>minut<br>e                                       |

# **Redis Server Metrics of DCS Redis Instances**

#### **NOTE**

- These metrics are supported for master/standby, read/write splitting, and cluster instances.
- **Dimensions** lists the metric dimensions.

| Table 14- | 4 Redis Server | <sup>-</sup> metrics |
|-----------|----------------|----------------------|
|-----------|----------------|----------------------|

| Metric ID         | Metri<br>c<br>Nam<br>e          | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-------------------|---------------------------------|---|--------------------|----------|--------------------------------|---|---|
| cpu_usag<br>e     | CPU<br>Usag<br>e                | The monitored<br>object's maximum<br>CPU usage among<br>multiple sampling<br>values in a<br>monitoring period | 0-100              | %        | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>)                   | 1<br>minut<br>e                                       |
| cpu_avg_<br>usage | Avera<br>ge<br>CPU<br>Usag<br>e | The monitored<br>object's average<br>CPU usage of<br>multiple sampling<br>values in a<br>monitoring period    | 0-100              | %        | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| memory_<br>usage  | Mem<br>ory<br>Usag<br>e         | Memory<br>consumed by the<br>monitored object   | 0-100              | %        | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>)                   | 1<br>minut<br>e                                       |

| Metric ID                       | Metri<br>c<br>Nam<br>e                      | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------------------|---|--|--------------------|----------|--------------------------------|---|---|
| connecte<br>d_clients           | Conn<br>ected<br>Client<br>s                | Number of<br>connected clients.<br>Includes<br>connections<br>established for<br>system<br>monitoring,<br>configuration<br>synchronization,<br>and services.<br>Excludes<br>connections from<br>replicas | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>)                 | 1<br>minut<br>e                                       |
| client_lon<br>gest_out_l<br>ist | Client<br>Long<br>est<br>Outp<br>ut<br>List | Longest output<br>list among current<br>client connections   | ≥ 0                | N/A      | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 3.0<br>and 4.0<br>instances<br>) | 1<br>minut<br>e                                       |
| client_big<br>gest_in_b<br>uf   | Client<br>Bigge<br>st<br>Input<br>Buf       | Maximum input<br>data length<br>among current<br>client connections  | ≥ 0                | byte     | 102<br>4(IE<br>C)              | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 3.0<br>and 4.0<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID           | Metri<br>c<br>Nam<br>e     | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------|----------------------------|---|--------------------|----------|--------------------------------|---|---|
| blocked_c<br>lients | Block<br>ed<br>Client<br>s | Number of clients<br>suspended by<br>block operations<br>such as BLPOP,<br>BRPOP, and<br>BRPOPLPUSH                               | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| used_me<br>mory     | Used<br>Mem<br>ory         | Total number of<br>bytes used by the<br>Redis server  | ≥ 0                | byte     | 102<br>4(IE<br>C)              | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| used_me<br>mory_rss | Used<br>Mem<br>ory<br>RSS  | RSS memory that<br>the Redis server<br>has used, which<br>includes all stack<br>and heap memory<br>but not swapped-<br>out memory | ≥ 0                | byte     | 102<br>4(IE<br>C)              | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID                | Metri<br>c<br>Nam<br>e                       | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|--------------------------|--|--|--------------------|----------|--------------------------------|---|---|
| used_me<br>mory_pea<br>k | Used<br>Mem<br>ory<br>Peak                   | Peak memory<br>consumed by<br>Redis since the<br>Redis server last<br>started                                | ≥ 0                | byte     | 102<br>4(IE<br>C)              | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| used_me<br>mory_lua      | Used<br>Mem<br>ory<br>Lua                    | Number of bytes<br>used by the Lua<br>engine   | ≥ 0                | byte     | 102<br>4(IE<br>C)              | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| memory_f<br>rag_ratio    | Mem<br>ory<br>Frag<br>ment<br>ation<br>Ratio | Current memory<br>fragmentation,<br>which is the ratio<br>between<br><b>used_memory_rs</b><br>s/used_memory. | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID                          | Metri<br>c<br>Nam<br>e            | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|------------------------------------|-----------------------------------|--|--------------------|----------|--------------------------------|---|---|
| total_con<br>nections_r<br>eceived | New<br>Conn<br>ectio<br>ns        | Number of<br>connections<br>received during<br>the monitoring<br>period. Includes<br>connections from<br>replicas and<br>established for<br>system<br>monitoring,<br>configuration<br>synchronization,<br>and services | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| total_com<br>mands_pr<br>ocessed   | Com<br>mand<br>s<br>Proce<br>ssed | Number of<br>commands<br>processed during<br>the monitoring<br>period  | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| instantan<br>eous_ops              | Ops<br>per<br>Secon<br>d          | Number of<br>commands<br>processed per<br>second   | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID                        | Metri<br>c<br>Nam<br>e             | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t  | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|----------------------------------|------------------------------------|--|--------------------|-----------|--------------------------------|---|---|
| total_net_<br>input_byt<br>es    | Netw<br>ork<br>Input<br>Bytes      | Number of bytes<br>received during<br>the monitoring<br>period | ≥ 0                | byte      | 102<br>4(IE<br>C)              | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| total_net_<br>output_by<br>tes   | Netw<br>ork<br>Outp<br>ut<br>Bytes | Number of bytes<br>sent during the<br>monitoring period        | ≥ 0                | byte      | 102<br>4(IE<br>C)              | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| instantan<br>eous_inpu<br>t_kbps | Input<br>Flow                      | Instantaneous<br>input traffic                                 | ≥ 0                | KiB/<br>s | 102<br>4(IE<br>C)              | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| Metric ID                         | Metri<br>c<br>Nam<br>e              | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t  | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------------------|-------------------------------------|--|--------------------|-----------|--------------------------------|---|---|
| instantan<br>eous_out<br>put_kbps | Outp<br>ut<br>Flow                  | Instantaneous<br>output traffic  | ≥ 0                | KiB/<br>s | 102<br>4(IE<br>C)              | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| rejected_c<br>onnection<br>s      | Reject<br>ed<br>Conn<br>ectio<br>ns | Number of<br>connections that<br>have exceeded<br>maxclients and<br>been rejected<br>during the<br>monitoring period | ≥ 0                | N/A       | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| expired_k<br>eys                  | Expire<br>d<br>Keys                 | Number of keys<br>that have expired<br>and been deleted<br>during the<br>monitoring period                           | ≥ 0                | N/A       | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID           | Metri<br>c<br>Nam<br>e     | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------|----------------------------|--|--------------------|----------|--------------------------------|---|---|
| evicted_k<br>eys    | Evicte<br>d<br>Keys        | Number of keys<br>that have been<br>evicted and<br>deleted during the<br>monitoring period | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| pubsub_c<br>hannels | PubS<br>ub<br>Chan<br>nels | Number of<br>Pub/Sub channels  | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| pubsub_p<br>atterns | PubS<br>ub<br>Patter<br>ns | Number of<br>Pub/Sub patterns  | ≥ 0                | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID                 | Metri<br>c<br>Nam<br>e                    | Metric<br>Description   | Value<br>Rang<br>e                         | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------------|---|---|--|----------|--------------------------------|---|---|
| keyspace_<br>hits_perc    | Hit<br>Rate                               | Ratio of the<br>number of Redis<br>cache hits to the<br>number of<br>lookups.<br>Calculation:<br>keyspace_hits/<br>(keyspace_hits +<br>keyspace_misses)<br>If no read<br>command is<br>performed within<br>a monitoring<br>period, the ratio is<br>0. | 0-100                                      | %        | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| command<br>_max_del<br>ay | Maxi<br>mum<br>Com<br>mand<br>Laten<br>Cy | Maximum latency<br>of commands  | ≥ 0  | ms       | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |
| is_slow_lo<br>g_exist     | Slow<br>Quer<br>y<br>Logs                 | Existence of slow<br>query logs in the<br>node<br>NOTE<br>Slow queries<br>caused by the<br>MIGRATE,<br>SLAVEOF,<br>CONFIG, BGSAVE,<br>and<br>BGREWRITEAOF<br>commands are not<br>counted.   | <ul> <li>1: ye s</li> <li>0: no</li> </ul> | N/A      | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID    | Metri<br>c<br>Nam<br>e | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|--------------|------------------------|--|--------------------|-------------|--------------------------------|---|---|
| keys         | Keys                   | Number of keys in<br>Redis                                   | ≥ 0                | N/A         | N/A                            | Data<br>node<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis<br>instances<br>)                   | 1<br>minut<br>e                                       |
| sadd         | SADD                   | Number of <b>SADD</b><br>commands<br>processed per<br>second | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| smember<br>s | SME<br>MBER<br>S       | Number of<br>SMEMBERS<br>commands<br>processed per<br>second | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID          | Metri<br>c<br>Nam<br>e | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)   | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|--------------------|------------------------|---|--------------------|-------------|--------------------------------|--|---|
| ms_repl_of<br>fset | Replic<br>ation<br>Gap | Data<br>synchronization<br>gap between the<br>master and the<br>replica | -                  | Byte        | 102<br>4(IE<br>C)              | Data<br>nodes<br>(standby<br>nodes of<br>master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| del                | DEL                    | Number of <b>DEL</b><br>commands<br>processed per<br>second             | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>)                        | 1<br>minut<br>e                                       |
| expire             | EXPIR<br>E             | Number of<br><b>EXPIRE</b><br>commands<br>processed per<br>second       | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>)                        | 1<br>minut<br>e                                       |

| Metric ID | Metri<br>c<br>Nam<br>e | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------|------------------------|--|--------------------|-------------|--------------------------------|---|---|
| get       | GET                    | Number of <b>GET</b><br>commands<br>processed per<br>second  | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| hdel      | HDEL                   | Number of <b>HDEL</b><br>commands<br>processed per<br>second | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| hget      | HGET                   | Number of <b>HGET</b><br>commands<br>processed per<br>second | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID | Metri<br>c<br>Nam<br>e | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------|------------------------|--|--------------------|-------------|--------------------------------|---|---|
| hmget     | HMG<br>ET              | Number of<br>HMGET<br>commands<br>processed per<br>second    | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| hmset     | HMS<br>ET              | Number of<br>HMSET<br>commands<br>processed per<br>second    | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>)                                  | 1<br>minut<br>e                                       |
| hset      | HSET                   | Number of <b>HSET</b><br>commands<br>processed per<br>second | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID | Metri<br>c<br>Nam<br>e | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------|------------------------|--|--------------------|-------------|--------------------------------|---|---|
| mget      | MGET                   | Number of <b>MGET</b><br>commands<br>processed per<br>second | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| mset      | MSET                   | Number of <b>MSET</b><br>commands<br>processed per<br>second | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| set       | SET                    | Number of <b>SET</b><br>commands<br>processed per<br>second  | 0–<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID           | Metri<br>c<br>Nam<br>e       | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t  | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------|------------------------------|--|--------------------|-----------|--------------------------------|---|---|
| rx_control<br>led   | Flow<br>Contr<br>ol<br>Times | Number of times<br>that client<br>requests are<br>controlled in a<br>period. This<br>metric is<br>incremented by 1<br>each time a client<br>request is<br>controlled.<br>If the value is<br>greater than 0,<br>the consumed<br>bandwidth<br>exceeds the upper<br>limit and flow<br>control is<br>triggered on a<br>node. The node<br>suspends client<br>commands<br>temporarily. | ≥ 0                | Cou<br>nt | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| bandwidt<br>h_usage | Band<br>width<br>Usag<br>e   | Percentage of the<br>used bandwidth<br>to the maximum<br>bandwidth limit   | 0-200              | %         | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID             | Metri<br>c<br>Nam<br>e          | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------|---------------------------------|---|--------------------|----------|--------------------------------|---|---|
| connectio<br>ns_usage | Conn<br>ectio<br>n<br>Usag<br>e | Percentage of the<br>current number of<br>connections to<br>the maximum<br>allowed number<br>of connections | 0-100              | %        | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| command<br>_max_rt    | Maxi<br>mum<br>Laten<br>cy      | Maximum delay<br>from when the<br>node receives<br>commands to<br>when it responds                          | ≥ 0                | μs       | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| command<br>_avg_rt    | Avera<br>ge<br>Laten<br>cy      | Average delay<br>from when the<br>node receives<br>commands to<br>when it responds                          | ≥ 0                | μs       | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID           | Metri<br>c<br>Nam<br>e | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------|------------------------|---|--------------------|-------------|--------------------------------|---|---|
| sync_full           | Full<br>Sync<br>Times  | Total number of<br>full<br>synchronizations<br>since the Redis<br>Server last started | ≥ 0                | N/A         | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| slow_log_<br>counts | Slow<br>Queri<br>es    | Number of times<br>that slow queries<br>occur within a<br>monitoring period           | ≥ 0                | cou<br>nt   | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| scan                | SCAN                   | Number of SCAN<br>operations per<br>second  | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |

| Metric ID         | Metri<br>c<br>Nam<br>e       | Metric<br>Description                       | Value<br>Rang<br>e | Uni<br>t    | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion)  | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-------------------|------------------------------|---|--------------------|-------------|--------------------------------|---|---|
| setex             | SETE<br>X                    | Number of SETEX<br>operations per<br>second | 0-<br>500,0<br>00  | Cou<br>nt/s | N/A                            | Data<br>nodes<br>(master/<br>standby,<br>read/<br>write<br>splitting,<br>and<br>cluster<br>Redis 4.0<br>and later<br>instances<br>) | 1<br>minut<br>e                                       |
| used_stor<br>age  | Used<br>Stora<br>ge<br>Space | Storage space<br>that has been<br>used.     | ≥ 0                | byte        | 102<br>4(IE<br>C)              | Data<br>nodes<br>(Redis<br>6.0<br>enterpris<br>e<br>(storage)<br>edition<br>instance)   | 1<br>minut<br>e                                       |
| storage_u<br>sage | Stora<br>ge<br>Usag<br>e     | Percentage of<br>used storage.              | 0-100              | %           | N/A                            | Data<br>nodes<br>(Redis<br>6.0<br>enterpris<br>e<br>(storage)<br>edition<br>instance)   | 1<br>minut<br>e                                       |

### **Proxy Metrics**

**NOTE** 

- These metrics are supported by Proxy Cluster and read/write splitting instances.
- **Dimensions** lists the metric dimensions.

| Metric ID               | Metri<br>c<br>Nam<br>e                                    | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t         | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-------------------------|---|---|--------------------|------------------|--------------------------------|--|---|
| cpu_usag<br>e           | CPU<br>Usag<br>e  | The monitored<br>object's maximum<br>CPU usage among<br>multiple sampling<br>values in a<br>monitoring period         | 0-100              | %                | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| memory_<br>usage        | Mem<br>ory<br>Usag<br>e                                   | Memory<br>consumed by the<br>monitored object   | 0-100              | %                | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| p_connect<br>ed_clients | Conn<br>ected<br>Client<br>s                              | Number of connected clients   | ≥ 0                | N/A              | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| max_rxpc<br>k_per_sec   | Max.<br>NIC<br>Data<br>Packe<br>t<br>Recei<br>ve<br>Rate  | Maximum<br>number of data<br>packets received<br>by the proxy NIC<br>per second during<br>the monitoring<br>period    | 0-100<br>0000<br>0 | Pac<br>ket/<br>s | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| max_txpc<br>k_per_sec   | Max.<br>NIC<br>Data<br>Packe<br>t<br>Trans<br>mit<br>Rate | Maximum<br>number of data<br>packets<br>transmitted by the<br>proxy NIC per<br>second during the<br>monitoring period | 0-100<br>0000<br>0 | Pac<br>ket/<br>s | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| max_rxkB<br>_per_sec    | Maxi<br>mum<br>Inbou<br>nd<br>Band<br>width               | Largest volume of<br>data received by<br>the proxy NIC per<br>second  | >= 0               | KiB/<br>s        | 102<br>4(IE<br>C)              | Proxies                                    | 1<br>minut<br>e                                       |

 Table 14-5 Proxy metrics of Proxy Cluster DCS Redis 3.0 instances

| Metric ID             | Metri<br>c<br>Nam<br>e   | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t         | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------|--|--|--------------------|------------------|--------------------------------|--|---|
| max_txkB<br>_per_sec  | Maxi<br>mum<br>Outb<br>ound<br>Band<br>width                     | Largest volume of<br>data transmitted<br>by the proxy NIC<br>per second  | >= 0               | KiB/<br>s        | 102<br>4(IE<br>C)              | Proxies                                    | 1<br>minut<br>e                                       |
| avg_rxpck<br>_per_sec | Avera<br>ge<br>NIC<br>Data<br>Packe<br>t<br>Recei<br>ve<br>Rate  | Average number<br>of data packets<br>received by the<br>proxy NIC per<br>second during the<br>monitoring period    | 0-100<br>0000<br>0 | Pac<br>ket/<br>s | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| avg_txpck<br>_per_sec | Avera<br>ge<br>NIC<br>Data<br>Packe<br>t<br>Trans<br>mit<br>Rate | Average number<br>of data packets<br>transmitted by the<br>proxy NIC per<br>second during the<br>monitoring period | 0-100<br>0000<br>0 | Pac<br>ket/<br>s | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| avg_rxkB_<br>per_sec  | Avera<br>ge<br>Inbou<br>nd<br>Band<br>width                      | Average volume<br>of data received<br>by the proxy NIC<br>per second   | >= 0               | KiB/<br>s        | 102<br>4(IE<br>C)              | Proxies                                    | 1<br>minut<br>e                                       |
| avg_txkB_<br>per_sec  | Avera<br>ge<br>Outb<br>ound<br>Band<br>width                     | Average volume<br>of data<br>transmitted by the<br>proxy NIC per<br>second   | >= 0               | KiB/<br>s        | 102<br>4(IE<br>C)              | Proxies                                    | 1<br>minut<br>e                                       |

| Table 14-6 Proxy metrics of Proxy Cluster or read/write splitting DCS Redis |
|---|
| 4.0/5.0/6.0 instances   |

| Metric ID             | Metri<br>c<br>Nam<br>e          | Metric<br>Description  | Value<br>Rang<br>e  | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------|---------------------------------|--|---|----------|--------------------------------|--|---|
| node_stat<br>us       | Proxy<br>Statu<br>s             | Indication of<br>whether the proxy<br>is normal.   | <ul> <li>0:<br/>No<br/>rm<br/>al</li> <li>1:<br/>Ab<br/>no<br/>rm<br/>al</li> </ul> | N/A      | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| cpu_usag<br>e         | CPU<br>Usag<br>e                | The monitored<br>object's maximum<br>CPU usage among<br>multiple sampling<br>values in a<br>monitoring period  | 0-100   | %        | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| cpu_avg_<br>usage     | Avera<br>ge<br>CPU<br>Usag<br>e | The monitored<br>object's average<br>CPU usage of<br>multiple sampling<br>values in a<br>monitoring period   | 0-100   | %        | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| memory_<br>usage      | Mem<br>ory<br>Usag<br>e         | Memory<br>consumed by the<br>monitored object  | 0-100   | %        | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| connecte<br>d_clients | Conn<br>ected<br>Client<br>s    | Number of<br>connected clients.<br>Includes<br>connections<br>established for<br>system<br>monitoring,<br>configuration<br>synchronization,<br>and services.<br>Excludes<br>connections from<br>replicas | ≥ 0   | N/A      | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |

| Metric ID                         | Metri<br>c<br>Nam<br>e             | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t  | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-----------------------------------|------------------------------------|---|--------------------|-----------|--------------------------------|--|---|
| instantan<br>eous_ops             | Ops<br>per<br>Secon<br>d           | Number of<br>commands<br>processed per<br>second  | ≥ 0                | N/A       | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| instantan<br>eous_inpu<br>t_kbps  | Input<br>Flow                      | Instantaneous<br>input traffic  | ≥ 0                | KiB/<br>s | 102<br>4(IE<br>C)              | Proxies                                    | 1<br>minut<br>e                                       |
| instantan<br>eous_out<br>put_kbps | Outp<br>ut<br>Flow                 | Instantaneous<br>output traffic   | ≥ 0                | KiB/<br>s | 102<br>4(IE<br>C)              | Proxies                                    | 1<br>minut<br>e                                       |
| total_net_<br>input_byt<br>es     | Netw<br>ork<br>Input<br>Bytes      | Number of bytes<br>received during<br>the monitoring<br>period  | ≥ 0                | byte      | 102<br>4(IE<br>C)              | Proxies                                    | 1<br>minut<br>e                                       |
| total_net_<br>output_by<br>tes    | Netw<br>ork<br>Outp<br>ut<br>Bytes | Number of bytes<br>sent during the<br>monitoring period   | ≥ 0                | byte      | 102<br>4(IE<br>C)              | Proxies                                    | 1<br>minut<br>e                                       |
| connectio<br>ns_usage             | Conn<br>ectio<br>n<br>Usag<br>e    | Percentage of the<br>current number of<br>connections to<br>the maximum<br>allowed number<br>of connections | 0-100              | %         | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| command<br>_max_rt                | Maxi<br>mum<br>Laten<br>cy         | Maximum delay<br>from when the<br>node receives<br>commands to<br>when it responds                          | ≥ 0                | μs        | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |
| command<br>_avg_rt                | Avera<br>ge<br>Laten<br>cy         | Average delay<br>from when the<br>node receives<br>commands to<br>when it responds                          | ≥ 0                | μs        | N/A                            | Proxies                                    | 1<br>minut<br>e                                       |

### **DCS Memcached Instance Metrics**

#### 

**Dimensions** lists the metric dimensions.

|  | Table | 14-7 | DCS | Memcached | instance | metrics |
|--|-------|------|-----|-----------|----------|---------|
|--|-------|------|-----|-----------|----------|---------|

| Metric ID                    | Metri<br>c<br>Nam<br>e                      | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t   | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|------------------------------|---|---|--------------------|------------|--------------------------------|--|---|
| cpu_usag<br>e                | CPU<br>Usag<br>e                            | The monitored<br>object's maximum<br>CPU usage among<br>multiple sampling<br>values in a<br>monitoring period | 0-100              | %          | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| memory_<br>usage             | Mem<br>ory<br>Usag<br>e                     | Memory<br>consumed by the<br>monitored object   | 0-100              | %          | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| net_in_thr<br>oughput        | Netw<br>ork<br>Input<br>Throu<br>ghput      | Inbound<br>throughput per<br>second on a port   | >= 0               | byte<br>/s | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| net_out_t<br>hroughpu<br>t   | Netw<br>ork<br>Outp<br>ut<br>Throu<br>ghput | Outbound<br>throughput per<br>second on a port  | >= 0               | byte<br>/s | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_conne<br>cted_clien<br>ts | Conn<br>ected<br>Client<br>s                | Number of<br>connected clients<br>(excluding those<br>from slave nodes)                                       | ≥ 0                | N/A        | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_used_<br>memory           | Used<br>Mem<br>ory                          | Number of bytes<br>used by<br>Memcached   | ≥ 0                | byte       | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |

| Metric ID                       | Metri<br>c<br>Nam<br>e                       | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|---------------------------------|--|--|--------------------|----------|--------------------------------|--|---|
| mc_used_<br>memory_r<br>ss      | Used<br>Mem<br>ory<br>RSS                    | RSS memory used<br>that actually<br>resides in the<br>memory, including<br>all stack and heap<br>memory but not<br>swapped-out<br>memory | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_used_<br>memory_<br>peak     | Used<br>Mem<br>ory<br>Peak                   | Peak memory<br>consumed since<br>the server last<br>started  | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_mem<br>ory_frag_r<br>atio    | Mem<br>ory<br>Frag<br>ment<br>ation<br>Ratio | Ratio between<br>Used Memory RSS<br>and Used Memory  | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_conne<br>ctions_rec<br>eived | New<br>Conn<br>ectio<br>ns                   | Number of<br>connections<br>received during<br>the monitoring<br>period  | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_com<br>mands_pr<br>ocessed   | Com<br>mand<br>s<br>Proce<br>ssed            | Number of<br>commands<br>processed during<br>the monitoring<br>period  | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_instan<br>taneous_<br>ops    | Ops<br>per<br>Secon<br>d                     | Number of<br>commands<br>processed per<br>second   | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_net_in<br>put_bytes          | Netw<br>ork<br>Input<br>Bytes                | Number of bytes<br>received during<br>the monitoring<br>period   | ≥ 0                | byte     | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |

| Metric ID                                | Metri<br>c<br>Nam<br>e                         | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t  | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|--|--|--|--------------------|-----------|--------------------------------|--|---|
| mc_net_o<br>utput_byt<br>es              | Netw<br>ork<br>Outp<br>ut<br>Bytes             | Number of bytes<br>sent during the<br>monitoring period  | ≥ 0                | byte      | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_instan<br>taneous_i<br>nput_kbp<br>s  | Input<br>Flow                                  | Instantaneous<br>input traffic   | ≥ 0                | KiB/<br>s | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_instan<br>taneous_<br>output_kb<br>ps | Outp<br>ut<br>Flow                             | Instantaneous<br>output traffic  | ≥ 0                | KiB/<br>s | 102<br>4(IE<br>C)              | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_reject<br>ed_conne<br>ctions          | Reject<br>ed<br>Conn<br>ectio<br>ns            | Number of<br>connections that<br>have exceeded<br>maxclients and<br>been rejected<br>during the<br>monitoring period | ≥ 0                | N/A       | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_expire<br>d_keys                      | Expire<br>d<br>Keys                            | Number of keys<br>that have expired<br>and been deleted<br>during the<br>monitoring period                           | ≥ 0                | N/A       | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_evicte<br>d_keys                      | Evicte<br>d<br>Keys                            | Number of keys<br>that have been<br>evicted and<br>deleted during the<br>monitoring period                           | ≥ 0                | N/A       | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_cmd_<br>get                           | Num<br>ber of<br>Retrie<br>val<br>Requ<br>ests | Number of<br>received data<br>retrieval requests   | ≥ 0                | N/A       | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |

| Metric ID            | Metri<br>c<br>Nam<br>e                       | Metric<br>Description   | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|----------------------|--|---|--------------------|----------|--------------------------------|--|---|
| mc_cmd_s<br>et       | Num<br>ber of<br>Stora<br>ge<br>Requ<br>ests | Number of<br>received data<br>storage requests                                    | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_cmd_f<br>lush     | Num<br>ber of<br>Flush<br>Requ<br>ests       | Number of<br>received data<br>clearance requests                                  | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_cmd_t<br>ouch     | Num<br>ber of<br>Touch<br>Requ<br>ests       | Number of<br>received requests<br>for modifying the<br>validity period of<br>data | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_get_hi<br>ts      | Num<br>ber of<br>Retrie<br>val<br>Hits       | Number of<br>successful data<br>retrieval<br>operations                           | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_get_<br>misses    | Num<br>ber of<br>Retrie<br>val<br>Misse<br>s | Number of failed<br>data retrieval<br>operations due to<br>key nonexistence       | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_delete<br>_hits   | Num<br>ber of<br>Delet<br>e Hits             | Number of<br>successful data<br>deletion<br>operations                            | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_delete<br>_misses | Num<br>ber of<br>Delet<br>e<br>Misse<br>s    | Number of failed<br>data deletion<br>operations due to<br>key nonexistence        | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |

| Metric ID          | Metri<br>c<br>Nam<br>e                                | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|--------------------|---|--|--------------------|----------|--------------------------------|--|---|
| mc_incr_h<br>its   | Num<br>ber of<br>Incre<br>ment<br>Hits                | Number of<br>successful<br>increment<br>operations                     | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_incr_<br>misses | Num<br>ber of<br>Incre<br>ment<br>Misse<br>s          | Number of failed<br>increment<br>operations due to<br>key nonexistence | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_decr_<br>hits   | Num<br>ber of<br>Decre<br>ment<br>Hits                | Number of<br>successful<br>decrement<br>operations                     | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_decr_<br>misses | Num<br>ber of<br>Decre<br>ment<br>Misse<br>s          | Number of failed<br>decrement<br>operations due to<br>key nonexistence | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_cas_hi<br>ts    | Num<br>ber of<br>CAS<br>Hits                          | Number of<br>successful CAS<br>operations                              | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_cas_m<br>isses  | Num<br>ber of<br>CAS<br>Misse<br>s                    | Number of failed<br>CAS operations<br>due to key<br>nonexistence       | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_cas_b<br>adval  | Num<br>ber of<br>CAS<br>Value<br>s Not<br>Matc<br>hed | Number of failed<br>CAS operations<br>due to CAS value<br>mismatch     | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |

| Metric ID                    | Metri<br>c<br>Nam<br>e                    | Metric<br>Description  | Value<br>Rang<br>e | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|------------------------------|---|--|--------------------|----------|--------------------------------|--|---|
| mc_touch<br>_hits            | Num<br>ber of<br>Touch<br>Hits            | Number of<br>successful<br>requests for<br>modifying the<br>validity period of<br>data                     | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_touch<br>_misses          | Num<br>ber of<br>Touch<br>Misse<br>s      | Number of failed<br>requests for<br>modifying the<br>validity period of<br>data due to key<br>nonexistence | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_auth_<br>cmds             | Authe<br>nticat<br>ion<br>Requ<br>ests    | Number of<br>authentication<br>requests  | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_auth_<br>errors           | Authe<br>nticat<br>ion<br>Failur<br>es    | Number of failed<br>authentication<br>requests   | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_curr_i<br>tems            | Num<br>ber of<br>Items<br>Store<br>d      | Number of stored<br>data items   | ≥ 0                | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_com<br>mand_ma<br>x_delay | Maxi<br>mum<br>Com<br>mand<br>Laten<br>cy | Maximum latency<br>of commands   | ≥ 0                | ms       | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |

| Metric ID                     | Metri<br>c<br>Nam<br>e    | Metric<br>Description   | Value<br>Rang<br>e                         | Uni<br>t | Con<br>vers<br>ion<br>Rul<br>e | Monitor<br>ed<br>Object<br>(Dimens<br>ion) | Moni<br>torin<br>g<br>Perio<br>d<br>(Raw<br>Data<br>) |
|-------------------------------|---------------------------|---|--|----------|--------------------------------|--|---|
| mc_is_slo<br>w_log_exi<br>st  | Slow<br>Quer<br>y<br>Logs | Existence of slow<br>query logs in the<br>instance<br>NOTE<br>Slow queries<br>caused by the<br>MIGRATE,<br>SLAVEOF,<br>CONFIG, BGSAVE,<br>and<br>BGREWRITEAOF<br>commands are not<br>counted. | <ul> <li>1: ye s</li> <li>0: no</li> </ul> | N/A      | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |
| mc_keysp<br>ace_hits_<br>perc | Hit<br>Rate               | Ratio of the<br>number of<br>Memcached cache<br>hits to the<br>number of<br>lookups   | 0-100                                      | %        | N/A                            | DCS<br>Memcac<br>hed<br>instance           | 1<br>minut<br>e                                       |

#### Dimensions

| Кеу                       | Value                         |
|---------------------------|-------------------------------|
| dcs_instance_id           | DCS Redis instance            |
| dcs_cluster_redis_node    | Redis Server                  |
| dcs_cluster_proxy_node    | Proxies (Redis 3.0)           |
| dcs_cluster_proxy2_node   | Proxies (Redis 4.0 and later) |
| dcs_memcached_instance_id | DCS Memcached instance        |

# **14.2 Common DCS Metrics**

This section describes common Redis metrics.

| Table 14- | 8 Common | metrics |
|-----------|----------|---------|
|-----------|----------|---------|

| Metric            | Description   |
|-------------------|---|
| CPU Usage         | This metric indicates the maximum value in each measurement period (minute-level: every minute; second-level: every 5 seconds).   |
|                   | <ul> <li>For a single-node or master/standby instance, you<br/>can view the CPU usage of the instance.</li> </ul>   |
|                   | <ul> <li>For a Proxy Cluster instance, you can view the CPU<br/>usage of the Redis Servers and the proxies.</li> </ul>  |
|                   | • For a Redis Cluster instance, you can only view the CPU usage of the Redis Servers.   |
| Memory Usage      | This metric measures the memory usage in each measurement period (minute-level: every minute; second-level: every 5 seconds).   |
|                   | <ul> <li>For a single-node or master/standby instance, you<br/>can view the memory usage of the instance.</li> </ul>  |
|                   | <ul> <li>For a Proxy Cluster instance, you can view the<br/>memory usage of the instance and the proxies.</li> </ul>  |
|                   | • For a Redis Cluster instance, you can only view the memory usage of the Redis Servers.  |
|                   | <b>NOTICE</b><br>The memory usage does not include the usage of reserved<br>memory.   |
| Connected Clients | This metric indicates the number of instantaneous connected clients, that is, the number of concurrent connections.   |
|                   | This metric does not include the number of connections to the standby nodes of master/standby or cluster instances.   |
|                   | For details about the maximum allowed number of connections, see the "Max. Connections" column of different instance types listed in DCS Instance Specifications.                             |
| Ops per Second    | This metric indicates the number of operations processed per second.  |
|                   | For details about the maximum allowed number of operations per second, see the "Reference Performance (QPS)" column of different instance types listed in <b>DCS</b> Instance Specifications. |
| Input Flow        | This metric indicates the instantaneous input traffic.  |
|                   | • The monitoring data on the instance level shows the aggregated input traffic of all nodes.  |
|                   | • The monitoring data on the node level shows the input traffic of the current node.  |

| Metric             | Description   |
|--------------------|---|
| Output Flow        | <ul> <li>This metric indicates the instantaneous output traffic.</li> <li>The monitoring data on the instance level shows the aggregated output traffic of all nodes.</li> <li>The monitoring data on the node level shows the output traffic of the current node.</li> </ul>   |
| Bandwidth Usage    | This metric indicates the percentage of the used<br>bandwidth to the maximum bandwidth limit.<br>Bandwidth usage = (Input flow + Output flow)/(2 ×<br>Maximum bandwidth) × 100%   |
| Commands Processed | This metric indicates the number of commands<br>processed during the monitoring period. The default<br>monitoring period is 1 minute.<br>The monitoring period of this metric is different from<br>that of the <b>Ops per Second</b> metric The <b>Ops per</b><br><b>Second</b> metric measures the instantaneous number of<br>commands processed. The <b>Commands Processed</b><br>metric measures the total number of commands<br>processed during the monitoring period. |
| Flow Control Times | This metric indicates the number of times that the<br>maximum allowed bandwidth is exceeded during the<br>monitoring period.<br>For details about the maximum allowed bandwidth, see<br>the "Maximum/Assured Bandwidth" column of different<br>instance types listed in DCS Instance Specifications.  |
| Slow Queries       | This metric indicates whether slow queries exist on the instance.<br>For details about the cause of a slow query, see<br>Viewing Redis Slow Queries.  |

# **14.3 Viewing DCS Metrics**

The Cloud Eye service monitors the running performance your DCS instances.

### Procedure

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

- **Step 2** Click O in the upper left corner of the console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** Click the desired instance.

# **Step 5** Choose **Performance Monitoring**. All monitoring metrics of the instance are displayed.

**NOTE** 

You can also click **View Metric** in the **Operation** column on the **Cache Manager** page. You will be redirected to the Cloud Eye console. The metrics displayed on the Cloud Eye console are the same as those displayed on the **Performance Monitoring** page of the DCS console.

----End

# 14.4 Configuring DCS Monitoring and Alarms

This section describes the alarm rules of some metrics and how to configure the rules. In actual scenarios, configure alarm rules for metrics by referring to the following alarm policies.

### **Alarm Policies for DCS Redis Instances**

| Metric       | Value<br>Range | Alarm<br>Policy   | Appro<br>ach<br>Upper<br>Limit | Handling Suggestion   |
|--------------|----------------|---|--------------------------------|---|
| CPU<br>Usage | 0-100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | No                             | Consider capacity expansion<br>based on the service analysis.<br>The CPU capacity of a single-<br>node or master/standby instance<br>cannot be expanded. If you need<br>larger capacity, use a cluster<br>instance instead.<br>This metric is available only for<br>single-node, master/standby, and<br>Proxy Cluster instances. For Redis<br>Cluster instances, this metric is<br>available only on the Redis<br>Server level. You can view the<br>metric on the <b>Redis Server</b> tab |
|              |                |   |                                | <b>Monitoring</b> page of the instance.   |

Table 14-9 DCS Redis instance metrics to configure alarm rules for

| Metric                   | Value<br>Range | Alarm<br>Policy  | Appro<br>ach<br>Upper<br>Limit | Handling Suggestion  |
|--------------------------|----------------|--|--------------------------------|--|
| Average<br>CPU<br>Usage  | 0-100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major    | No                             | Consider capacity expansion<br>based on the service analysis.<br>The CPU capacity of a single-<br>node or master/standby instance<br>cannot be expanded. If you need<br>larger capacity, use a cluster<br>instance instead.<br>This metric is available only for<br>single-node, master/standby, and<br>Proxy Cluster instances. For Redis<br>Cluster instances, this metric is<br>available only on the Redis<br>Server level. You can view the<br>metric on the <b>Redis Server</b> tab<br>page on the <b>Performance</b><br><b>Monitoring</b> page of the instance. |
| Memory<br>Usage          | 0-100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Critical | No                             | Expand the capacity of the instance.   |
| Connect<br>ed<br>Clients | 0-10,000       | Alarm<br>threshold:<br>> 8000<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major   | No                             | Optimize the connection pool in<br>the service code to prevent the<br>number of connections from<br>exceeding the maximum limit.<br>Configure this alarm policy on<br>the instance level for single-node<br>and master/standby instances.<br>For cluster instances, configure<br>this alarm policy on the Redis<br>Server and Proxy level.<br>For single-node and master/<br>standby instances, the maximum<br>number of connections allowed is<br>10,000. You can adjust the<br>threshold based on service<br>requirements.   |

| Metric                                    | Value<br>Range | Alarm<br>Policy   | Appro<br>ach<br>Upper<br>Limit | Handling Suggestion  |
|---|----------------|---|--------------------------------|--|
| New<br>Connecti<br>ons<br>(Count/<br>min) | ≥ 0            | Alarm<br>threshold:<br>> 10,000<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Minor                                | -                              | Check whether <b>connect</b> is used<br>and whether the client<br>connection is abnormal. Use<br>persistent connections<br>(" <b>pconnect</b> " in Redis<br>terminology) to ensure<br>performance.<br>Configure this alarm policy on<br>the instance level for single-node<br>and master/standby instances.<br>For cluster instances, configure<br>this alarm policy on the Redis<br>Server and Proxy level. |
| Input<br>Flow                             | ≥ 0            | Alarm<br>threshold:<br>> 80% of<br>the<br>assured<br>bandwidth<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | Yes                            | Consider capacity expansion<br>based on the service analysis and<br>bandwidth limit.<br>Configure this alarm only for<br>single-node and master/standby<br>DCS Redis 3.0 instances and set<br>the alarm threshold to 80% of<br>the assured bandwidth of DCS<br>Redis 3.0 instances.  |
| Output<br>Flow                            | ≥ 0            | Alarm<br>threshold:<br>> 80% of<br>the<br>assured<br>bandwidth<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | Yes                            | Consider capacity expansion<br>based on the service analysis and<br>bandwidth limit.<br>Configure this alarm only for<br>single-node and master/standby<br>DCS Redis 3.0 instances and set<br>the alarm threshold to 80% of<br>the assured bandwidth of DCS<br>Redis 3.0 instances.  |

## **Alarm Policies for DCS Memcached Instances**

| Metric                   | Value<br>Range | Alarm<br>Policy  | Approac<br>h Upper<br>Limit | Handling Suggestion   |
|--------------------------|----------------|--|-----------------------------|---|
| CPU<br>Usage             | 0-100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major    | No                          | Check the service for traffic<br>surge.<br>The CPU capacity of a single-<br>node or master/standby<br>instance cannot be expanded.<br>Analyze the service and<br>consider splitting the service or<br>combine multiple instances<br>into a cluster on the client end. |
| Memory<br>Usage          | 0-100%         | Alarm<br>threshold:<br>> 65%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Minor    | No                          | Consider expanding the instance capacity.   |
| Connect<br>ed<br>Clients | 0-10,000       | Alarm<br>threshold:<br>> 8000<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major   | No                          | Optimize the connection pool<br>in the service code to prevent<br>the number of connections<br>from exceeding the maximum<br>limit.   |
| New<br>Connecti<br>ons   | ≥ 0            | Alarm<br>threshold:<br>> 10,000<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Minor | -                           | Check whether <b>connect</b> is used<br>and whether the client<br>connection is abnormal. Use<br>persistent connections<br>("pconnect" in Redis<br>terminology) to ensure<br>performance.   |

| Metric                         | Value<br>Range | Alarm<br>Policy   | Approac<br>h Upper<br>Limit | Handling Suggestion   |
|--------------------------------|----------------|---|-----------------------------|---|
| Input<br>Flow                  | ≥ 0            | Alarm<br>threshold:<br>> 80% of<br>the<br>assured<br>bandwidth<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | Yes                         | Consider capacity expansion<br>based on the service analysis<br>and bandwidth limit.<br>For details about the<br>bandwidth of different instance<br>specifications, see DCS<br>Instance Specifications. |
| Output<br>Flow                 | ≥ 0            | Alarm<br>threshold:<br>> 80% of<br>the<br>assured<br>bandwidth<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | Yes                         | Consider capacity expansion<br>based on the service analysis<br>and bandwidth limit.<br>For details about the<br>bandwidth of different instance<br>specifications, see DCS<br>Instance Specifications. |
| Authenti<br>cation<br>Failures | ≥ 0            | Alarm<br>threshold:<br>> 0<br>Number of<br>consecutiv<br>e periods:<br>1<br>Alarm<br>severity:<br>Critical                                  | -                           | Check whether the password is<br>entered correctly.   |

### Alarm Policies for Redis Server Nodes of DCS Redis Instances

| Metric                  | Value<br>Range | Alarm<br>Policy   | Approac<br>h Upper<br>Limit | Handling Suggestion  |
|-------------------------|----------------|---|-----------------------------|--|
| CPU<br>Usage            | 0–100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | No                          | Check the service for traffic<br>surge.<br>Check whether the CPU usage<br>is evenly distributed to Redis<br>Server nodes. If the CPU usage<br>is high on multiple nodes,<br>consider capacity expansion.<br>Expanding the capacity of a<br>cluster instance will scale out<br>nodes to share the CPU<br>pressure.<br>If the CPU usage is high on a<br>single node, check whether hot<br>keys exist. If yes, optimize the<br>service code to eliminate hot<br>keys. |
| Average<br>CPU<br>Usage | 0–100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | No                          | Consider capacity expansion<br>based on the service analysis.<br>The CPU capacity of a single-<br>node, read/write splitting, or<br>master/standby instance<br>cannot be expanded. If you<br>need larger capacity, use a<br>cluster instance instead.  |
| Memory<br>Usage         | 0-100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | No                          | Check the service for traffic<br>surge.<br>Check whether the memory<br>usage is evenly distributed to<br>Redis Server nodes. If the<br>memory usage is high on<br>multiple nodes, consider<br>capacity expansion. If the<br>memory usage is high on a<br>single node, check whether big<br>keys exist. If yes, optimize the<br>service code to eliminate big<br>keys.  |

Table 14-11 Redis server metrics to configure alarm policies for

| Metric                   | Value<br>Range | Alarm<br>Policy  | Approac<br>h Upper<br>Limit | Handling Suggestion  |
|--------------------------|----------------|--|-----------------------------|--|
| Connect<br>ed<br>Clients | 0–10,000       | Alarm<br>threshold:<br>> 8000<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major   | No                          | Check whether the number of<br>connections is within the<br>appropriate range. If yes, adjust<br>the alarm threshold.                  |
| New<br>Connecti<br>ons   | ≥ 0            | Alarm<br>threshold:<br>> 10,000<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Minor | -                           | Check whether <b>connect</b> is<br>used. To ensure performance,<br>use persistent connections<br>("pconnect" in Redis<br>terminology). |
| Slow<br>Query<br>Logs    | 0–1            | Alarm<br>threshold:<br>> 0<br>Number of<br>consecutiv<br>e periods:<br>1<br>Alarm<br>severity:<br>Major      | -                           | Use the slow query function on<br>the console to analyze slow<br>commands.   |

| Metric                   | Value<br>Range | Alarm<br>Policy  | Approac<br>h Upper<br>Limit | Handling Suggestion  |
|--------------------------|----------------|--|-----------------------------|--|
| Bandwid<br>th Usage      | 0-200%         | Alarm<br>threshold:<br>> 90%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major  | Yes                         | Check whether the bandwidth<br>usage increase comes from<br>read services or write services<br>based on the input and output<br>flow.<br>If the bandwidth usage of a<br>single node is high, check<br>whether big keys exist.<br>Even if the bandwidth usage<br>exceeds 100%, flow control<br>may not necessarily be<br>performed. The actual flow<br>control is subject to the <b>Flow</b><br><b>Control Times</b> metric.<br>Even if the bandwidth usage is<br>below 100%, flow control may<br>be performed. The real-time<br>bandwidth usage is reported<br>once in every reporting period.<br>The flow control times metric<br>is reported every second.<br>During a reporting period, the<br>traffic may surge within<br>seconds and then fall back. By<br>the time the bandwidth usage<br>is reported, it has restored to<br>the normal level. |
| Flow<br>Control<br>Times | ≥ 0            | Alarm<br>threshold:<br>> 0<br>Number of<br>consecutiv<br>e periods:<br>1<br>Alarm<br>severity:<br>Critical | Yes                         | Consider capacity expansion<br>based on the specification<br>limits, input flow, and output<br>flow.<br><b>NOTE</b><br>This metric is supported only by<br>Redis 4.0 and later and not by<br>Redis 3.0.  |

### Alarm Policies for Proxy Nodes of DCS Redis Instances

| Metric                   | Value<br>Range | Alarm<br>Policy  | Approac<br>h Upper<br>Limit | Handling Suggestion   |
|--------------------------|----------------|--|-----------------------------|---|
| CPU<br>Usage             | 0–100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Critical | Yes                         | Consider capacity expansion, which will add proxies.  |
| Memory<br>Usage          | 0-100%         | Alarm<br>threshold:<br>> 70%<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Critical | Yes                         | Consider capacity expansion,<br>which will add proxies.   |
| Connect<br>ed<br>Clients | 0-30,000       | Alarm<br>threshold:<br>> 20,000<br>Number of<br>consecutiv<br>e periods:<br>2<br>Alarm<br>severity:<br>Major | No                          | Optimize the connection pool<br>in the service code to prevent<br>the number of connections<br>from exceeding the maximum<br>limit. |

 Table 14-12 Proxy metrics to configure alarm policies for

### Configuring an Alarm Rule for a Resource Group

Cloud Eye allows you to add DCS instances, Redis Server nodes, and proxy nodes to resource groups and manage instances and alarm rules by group to simplify O&M. For details, see **Creating a Resource Group**.

- **Step 1** Create a resource group.
  - 1. Log in to the Cloud Eye console. In the navigation pane, choose **Resource Groups** and then click **Create Resource Group** in the upper right corner.

Enter a group name and add Redis Server nodes to the resource group.
 You can add Redis Server nodes of different instances to the same resource group.

| Figure 14-1 | Creating | а | resource | group |
|-------------|----------|---|----------|-------|
|-------------|----------|---|----------|-------|

| < Create Resource Group  |  |  |   |            |  |  |  |
|--|--|--|---|------------|--|--|--|
| Group Name<br>Enterprise Project<br>Associate Enterprise Project | redis-server<br>default<br>Later Now 3   | C Create Enterprise Project  |   |            |  |  |  |
| Select resource  | Elastic Cloud Server (0)<br>Bare Metal Server (0)<br>Elastic Volume Service (0)                      | All resources  |   | Name • Q C |  |  |  |
|  | Elastic IP and Bandwidth (0)     Relational Database Service (0)     Distributed Message Service (0) | Name     dcs-6   | ID<br>739817c2-6d93-41d8-9e96-f8572028b | 24         |  |  |  |
|  | Distributed Cache Service (2)     DCS Memcached Instanc  | 2         Select All           2         Automation 0.000 Million 0.000  |   |            |  |  |  |
| DCS Redis Instances (0)     Redis Servers (2)     Remote (0)     |  | Implementation         Implementation           Implementation <td< th=""></td<> |   |            |  |  |  |
|  | Elastic Load Balance (0)     Cloud Search Service (0)     GaussDB NoSOL (0)                          |  |   |            |  |  |  |
|  | Document Database Service (0)     GaussD8(for MySQL) (0)     GaussD8(for openGauss) (0)              |  |   |            |  |  |  |
| You have selected 2 resources for the                            | current group. (You can select 1998 more resour  | ces.)  |   | Create     |  |  |  |

- 3. Click Create.
- Step 2 In the navigation pane of the Cloud Eye console, choose Alarm Management > Alarm Rules and then click Create Alarm Rule to set alarm information for the resource group.

Create a CPU usage alarm rule for all Redis Server nodes in the resource group, as shown in the following figure.

| Create Alarm Rule  |   |                     |
|--------------------|---|---------------------|
| * Name             | alam-tīj  |                     |
| Description        | 0256  |                     |
| * Alarm Type       | Metric Event  |                     |
| * Resource Type    | Distributed Cache Service • (2)   |                     |
| * Dimension        | DCS Redis Instances - Redis Servers •   |                     |
| * Monitoring Scope | All resources Resource groups, alarms will be higgered as resources under that group reach their predefined thresholds.                       |                     |
| * Group            | failover   C Create Resource Group View Resource Details in a Group   |                     |
| * Method           | Associate template Use existing template Configure manually   |                     |
| * Alarm Policy     | Name Alarm Policy   | Alarm Severity Open |
|                    | JUsage         ▼         Raw data         >=         ▼         70         % [3 times (consecutively)         ▼         Then One day         ▼ | Major 💌             |

Figure 14-2 Creating an alarm rule for a resource group

Step 3 Click Create.

----End

### Configuring an Alarm Rule for a Specific Resource

In the following example, an alarm rule is set for the **Slow Query Logs** (**is\_slow\_log\_exist**) metric.

- **Step 1** Log in to the **DCS console**.
- **Step 2** Click O in the upper left corner of the management console and select the region where your instance is located.
- **Step 3** In the navigation pane, choose **Cache Manager**.
- **Step 4** In the row containing the DCS instance whose metrics you want to view, click **View Metric** in the **Operation** column.

Figure 14-3 Viewing instance metrics

| Name ↓≡                        | Status ↓Ξ | Cache Engine ↓Ξ | Туре          | CPU | Specificatio ↓Ξ | Used/Availa ↓Ξ | Connection Addr | Enterprise Project | Tags | Billing Mode ⑦ Operation                                   |
|--------------------------------|-----------|-----------------|---------------|-----|-----------------|----------------|-----------------|--------------------|------|--|
| □ dcs-lm 2<br>0c2333-3a0b-4be4 | Running   | Basic Redis 5.0 | Redis Cluster | x86 | 8               | 9/8,           | redis-0c25f6    | default            |      | Pay-per-use<br>Created on May View Metric Restart   More + |

Step 5 On the displayed page, locate the Slow Query Logs metric. Hover over the metric

and click is to create an alarm rule for the metric.

The Create Alarm Rule page is displayed.

- Step 6 Specify the alarm information.
  - 1. Set the alarm name and description.
  - 2. Specify the alarm policy and alarm severity.

For example, the alarm policy shown in **Figure 14-4** indicates that an alarm will be triggered if slow queries exist in the instance for two consecutive periods. If no actions are taken, the alarm will be triggered once every day, until the value of this metric returns to **0**.

Figure 14-4 Setting the alarm content

| * Method C                   | nfigure manually  |                          |
|------------------------------|---|--------------------------|
| * Alarm Policy               |   |                          |
| Metric Name                  | Alarm Policy  | Alarm Severity Operation |
| Slow Query Logs              | Raw d ▼         2 consecuti ▼         >         ▼         0         One day         ▼ | Major 🔻                  |
| Add Alarm Policy You can add | 0 more.   |                          |

- 3. Set the alarm notification configurations. If you enable **Alarm Notification**, set the validity period, notification object, and trigger condition.
- 4. Click **Create**.

**NOTE** 

- For more information about creating alarm rules, see Creating an Alarm Rule.
- To modify or disable alarms, see Alarm Rule Management.

----End
## **15** Viewing DCS Audit Logs

With CTS, you can query, audit, and review operations performed on cloud resources. Traces include the operation requests sent using the management console or open APIs as well as the results of these requests.

## **DCS Operations Supported by CTS**

| Operation  | Resource<br>Type | Trace Name                     |
|--|------------------|--------------------------------|
| Creating an instance                             | Redis            | createDCSInstance              |
| Submitting<br>an instance<br>creation<br>request | Redis            | submitCreateDCSInstanceRequest |
| Deleting<br>multiple<br>instances                | Redis            | batchDeleteDCSInstance         |
| Deleting an instance                             | Redis            | deleteDCSInstance              |
| Modifying<br>instance<br>information             | Redis            | modifyDCSInstanceInfo          |
| Modifying<br>instance<br>configuratio<br>ns      | Redis            | modifyDCSInstanceConfig        |
| Changing<br>instance<br>password                 | Redis            | modifyDCSInstancePassword      |

Table 15-1 DCS operations that can be recorded by CTS

| Operation  | Resource<br>Type | Trace Name                           |
|--|------------------|--------------------------------------|
| Stopping an instance   | Redis            | stopDCSInstance                      |
| Submitting<br>an instance<br>stopping<br>request                 | Redis            | submitStopDCSInstanceRequest         |
| Restarting<br>an instance  | Redis            | restartDCSInstance                   |
| Submitting<br>an instance<br>restarting<br>request               | Redis            | submitRestartDCSInstanceRequest      |
| Starting an instance   | Redis            | startDCSInstance                     |
| Submitting<br>an instance<br>starting<br>request                 | Redis            | submitStartDCSInstanceRequest        |
| Clearing<br>instance<br>data                                     | Redis            | flushDCSInstance                     |
| Stopping<br>multiple<br>instances                                | Redis            | batchStopDCSInstance                 |
| Submitting a<br>request to<br>stop<br>instances in<br>batches    | Instance         | submitBatchStopDCSInstanceRequest    |
| Restarting<br>instances in<br>batches                            | Redis            | batchRestartDCSInstance              |
| Submitting a<br>request to<br>restart<br>instances in<br>batches | Redis            | submitBatchRestartDCSInstanceRequest |
| Starting<br>multiple<br>instances                                | Redis            | batchStartDCSInstance                |

| Operation  | Resource<br>Type | Trace Name                         |
|--|------------------|------------------------------------|
| Submitting a<br>request to<br>start<br>instances in<br>batches         | Instance         | submitBatchStartDCSInstanceRequest |
| Restoring<br>instance<br>data  | Redis            | restoreDCSInstance                 |
| Submitting a<br>request to<br>restore<br>instance<br>data              | Redis            | submitRestoreDCSInstanceRequest    |
| Backing up<br>instance<br>data   | Redis            | backupDCSInstance                  |
| Submitting a<br>request to<br>back up<br>instance<br>data              | Redis            | submitBackupDCSInstanceRequest     |
| Deleting<br>instance<br>backup files                                   | Redis            | deleteInstanceBackupFile           |
| Deleting<br>background<br>tasks  | Redis            | deleteDCSInstanceJobRecord         |
| Modifying<br>instance<br>specification<br>s                            | Redis            | modifySpecification                |
| Submitting a<br>request to<br>modify<br>instance<br>specification<br>s | Redis            | submitModifySpecificationRequest   |
| Creating an<br>instance<br>subscription<br>order                       | Redis            | createInstanceOrder                |

| Operation   | Resource<br>Type | Trace Name                     |
|---|------------------|--------------------------------|
| Creating an<br>order for<br>modifying<br>instance<br>specification<br>s | Redis            | createSpecificationChangeOrder |
| Updating<br>enterprise<br>project ID                                    | Redis            | updateEnterpriseProjectId      |
| Switching<br>between<br>master and<br>standby<br>nodes                  | Redis            | masterStandbySwitchover        |
| Disabling<br>public access  | Redis            | disablePublicNetworkAccess     |
| Enabling<br>public access   | Redis            | enablePublicNetworkAccess      |
| Resetting<br>instance<br>password                                       | Redis            | resetDCSInstancePassword       |
| Submitting a<br>request to<br>clear<br>instance<br>data                 | Redis            | submitFlushDCSInstanceRequest  |
| Accessing<br>Web CLI  | Redis            | webCliLogin                    |
| Running<br>commands<br>in Web CLI                                       | Redis            | webCliCommand                  |
| Exiting Web<br>CLI  | Redis            | webCliLogout                   |
| Migrating offline data  | Redis            | offlineMigrate                 |
| Changing<br>the billing<br>mode   | Redis            | billingModeChange              |
| Updating<br>instance tags   | Redis            | updateInstanceTag              |

| Operation                                       | Resource<br>Type | Trace Name      |
|---|------------------|-----------------|
| Modifying<br>the whitelist<br>configuratio<br>n | Instance         | modifyWhiteList |
| Modifying<br>instance<br>bandwidth              | Redis            | modifyBandwidth |

## Viewing Audit Logs

View CTS logs of DCS, see **Querying Real-Time Traces**.