

TencentDB for Redis

Troubleshooting

Product Documentation



Copyright Notice

©2013-2024 Tencent Cloud. All rights reserved.

Copyright in this document is exclusively owned by Tencent Cloud. You must not reproduce, modify, copy or distribute in any way, in whole or in part, the contents of this document without Tencent Cloud's the prior written consent.

Trademark Notice



All trademarks associated with Tencent Cloud and its services are owned by Tencent Cloud Computing (Beijing) Company Limited and its affiliated companies. Trademarks of third parties referred to in this document are owned by their respective proprietors.

Service Statement

This document is intended to provide users with general information about Tencent Cloud's products and services only and does not form part of Tencent Cloud's terms and conditions. Tencent Cloud's products or services are subject to change. Specific products and services and the standards applicable to them are exclusively provided for in Tencent Cloud's applicable terms and conditions.

Contents

Troubleshooting

Connection Exception

- One-Click Connectivity Checker

- Private Network Unable to Connect Location Guide

Exception Analysis and Solution of Redisson Client Timeout Reconnection

Performance Troubleshooting and Fine-Tuning

- Redis Instance Connection Failure

- High CPU Utilization

- High Outbound Traffic

- High Memory Utilization

- High Number of Total Requests

- Execution Error

Troubleshooting

Connection Exception

One-Click Connectivity Checker

Last updated : 2024-03-29 15:48:23

If the connection to TencentDB for Redis fails, it is recommended that you use the One-Click Connectivity Checker to locate the cause. This article introduces the potential causes and solutions for connection failures using the connectivity checker.

When a private network connection fails, see [Private Network One-Click Connectivity Check](#) to quickly locate the issue. If the problem persists, further analyze the cause of the inability to connect by following the [Guide to Troubleshooting Private Network Connectivity Issues](#).

For public network connection failures, see [Public Network One-Click Connectivity Check](#).

Private Network One-Click Connectivity Check

The Private Network One-Click Connectivity Check will inspect the following:

Whether the Redis and CVM instances are normal.

Whether CVM and Redis belong to the same VPC.

Whether the inbound rules of the Redis security group allow the IP address of CVM, and whether the outbound rules of the CVM security group allow the private IP address of Redis.

Note:

In the **Classic Network**, the security group bound to the CVM **cannot filter** data packets to or from No SQL Database (Redis, Memcached). Use the One-Click Connectivity Checker for the Classic Network security group's configuration. If the check results are abnormal, please perform manual inspection and verification.

Directions

1. Log in to the [Redis console](#).
2. Select the instance you want to troubleshoot, click **Instance ID** to open the instance details page.
3. On the instance details page, go to the **Connection Check > Private Network Check** page.

Private Network Check Public Network Check

Instance ID Instance Name Private IP address Private Network Port Network

You haven't performed the connection status check yet.

The connection check will help you detect the potential connection and access issues in your cloud database, and give recommended solutions to ensure you can access your cloud database normally.

[Common Causes of Database Connection Failure](#)

Initiate Check

Add CVMs to Access This Instance

4. Click **Add CVMs to Access This Instance**, in the following displayed window, all CVMs in the same region as the instance will automatically be displayed. Select the CVM used to access this Redis instance and click **OK**.

Note:

By default, only cloud services in the same region are provided. Cloud services across regions do not support configuration checks.

Add CVMs to Access This Instance

Select CVM ⓘ

Search for names or private IP addresses.

ins-
Private IP

Selected

ins-rymzbueg
Private IP

ins-
Private IP

OK **Cancel**

5. Click **Start Check**, wait for the check to complete. Then a check report will be generated, as shown in the following figure.

Abnormal

You haven't performed the connection status check yet.

The connection check will help you detect the potential connection and access issues in your cloud database, and give recommended solutions to ensure you can access your cloud database normally.

[Common Causes of Database Connection Failure](#)

[Initiate Check](#) End time of the last check: 2024-03-13 15:03:31

[Add CVMs to Access This Instance](#)

CVM Name	Private IP address	Network	Status
ins- st			Failed

6. In the **Operation** column, click **Peek Report**. Based on the check report, locate the issues and make adjustments according to the recommended solutions before attempting to reconnect to Redis.

Check Report Details

CVM Name ins-n1gmmhfo

Check Item	Status	Impact	Pro
Redis instance status	Passed	--	
CVM Instance Status	Passed	--	
CVM and Redis are in the same VPC	Abnormal	CVM cannot access the Redis instance.	It h anc VP the Re Pl
CVM Security Group Policy	Passed	--	
Redis security group policy	Passed	--	

Total items: 5

Cancel

Exception Analysis and Measures

Check Items	Status	Impact	Suggestion	Solution
Redis instance status	Abnormal	Redis instance unreachable	Isolated: Your Redis instance has been isolated. If you need to continue	For specific operations, see Restoring Isolated Instance .

			using this Redis instance, please go to the Redis Recycle Bin to restore the isolated instance.	
CVM instance status	Abnormal	CVM instance unreachable	<p>Isolated: Your CVM instance has been isolated. If you need to continue using this CVM instance, please go to the CVM Recycle Bin to restore the instance.</p> <p>Shutdown: Your CVM instance has been shut down. If you need to continue using this CVM instance, please go to the CVM console to start the CVM instance.</p>	<p>Isolated: For specific operations, see Instance Repossession or Recovering.</p> <p>Shutdown: For specific operations, see Starting Up Instances.</p>
CVM and Redis are in the same VPC	Abnormal	Redis instance unreachable from the CVM	Your server and the Redis instance are detected to be in different VPC IP ranges. The CVM instance needs to be in the same region and the same VPC as the Redis instance.	<p>To modify the network information of Redis, see Configuring Network.</p> <p>To modify the network information of CVM, see Switching to VPC.</p>
Redis security group policy	Abnormal	Redis instance unreachable from the CVM	It has been detected that the inbound rules of the security group bound with your Redis have not allowed access to the CVM's internal network IP address and ports.	<ol style="list-style-type: none"> 1. Log in to the CVM console, in the instance list under the primary IPv4 address column, verify the internal network IP address of the CVM connecting to Redis. For specific operations, see Viewing Instance Information. 2. Please go to the Redis console, under the Security Group tab to the Rule Preview section, check Inbound Rules. For specific operations, see Configuring Security Group.
CVM security	Abnormal	Redis instance	It has been detected that the outbound rules of the security	1. Log in to the CVM console , on the Security

group policy		unreachable from the CVM	group affiliated with your CVM have not allowed access to Redis's internal network IP address and ports.	Group management page, verify the Outbound Rules . For specific operations, see Viewing Security Groups . 2. Modify the outbound rules of the security group to allow access to Redis's internal network IP address and ports. For specific operations, see Modifying Security Group Rules .
--------------	--	--------------------------	--	--

Public Network One-Click Connectivity Check

The Public Network One-Click Connectivity Checker will inspect the following:

Whether the status of the Redis instance is normal.

Whether the instance is enabled with the public network access feature.

Whether the Redis security group allows access to external network addresses and ports.


Directions

1. Log in to the [Redis console](#).
2. Select the instance you want to troubleshoot, click on the **Instance ID** to enter the instance details page.
3. On the instance details page, go to the **Connection Check > Public Network Check** page.
4. Click **Add the Public server Addresses accessing this instance**. In the **Public Server Addresses** input field, enter the required external network server address according to the instructions, and then click **OK**.

Note:

The required format for the added external network server are as follows:

1. IP address format: xx.xx.xx.xx. Other formats will result in an error prompt.
2. Multiple IP address separators: Please use line break, comma, semicolon, or space.
3. Quantity limit: Up to 20 external network server addresses can be added for each check.



5. Click **Start Check**, and wait for the check to complete. Then a check report will be generated.

6. In the **Operation** column, click **Peek Report**. Based on the check report, locate the issues and make adjustment according to the recommended solutions before attempting to reconnect to Redis.

Check Report Details			
Public Network Server Address 0.0.0.0			
Check Item	Status	Impact	Processing Sug
Redis instance status	Passed	--	
Public Network Access Status	Abnormal	Public network server cannot access the Redis instance.	It is detected that has not enabled Please refer here
Redis security group policy	Passed	--	
Total items: 3			
Cancel			

Exception Analysis and Measures

Check Items	Status	Impact	Suggestion	Solution
Redis instance status	Abnormal	Redis instance unreachable	Isolated: Your Redis instance has been isolated. If you need to continue using this Redis instance, please go to the Redis Recycle Bin	For specific operations, see Restoring Isolated Instance .

			to restore the isolated instance.	
Public network access status	Abnormal	Public network servers are unable to access the Redis instance.	It has been detected that your Redis instance's public network is not enabled. To enable it, see Configuring Public Network Address .	To enable the public network, please refer to Configuring Public Network Address . Note: Currently, only the Chengdu, Beijing, Shanghai, and Guangzhou regions support public network addresses.
Redis security group policy	Abnormal	Public network servers are unable to access the Redis instance.	It has been detected that the inbound rules of the security group bound to your Redis instance have not allowed access to 1.1.1.1 and TCP protocol port 6379. If you have not configured it purposely, it may prevent your public network servers from accessing the Redis instance properly.	For detailed operations, see Configuring Security Group . Note: After public network access is enabled, it will be governed by the security group network access policy. Please configure the source information for database access in the security group's inbound rules and enable the protocol ports (both private and public network ports need to be enabled, and the default private network port is 6379).

Private Network Unable to Connect Location Guide

Last updated : 2024-03-29 15:54:07

Issue Description

Issue 1: Use the CVM to connect to a TencentDB for Redis instance at the private network address automatically assigned by the system as instructed in [Connecting to TencentDB for Redis Instance](#). However, the connection fails.

Issue 2: Log in to the [Redis console](#), go to the instance list, click **Log In** in the **Operation** column of the target instance to redirect to DMC, connect to the instance. However, the connection fails, as shown in the following figure.

Possible Causes

For database connection failures for the first time, the possible causes are as follows:

Port error.

Network configuration error or incorrect security group configuration.

Password error.

For sudden connection failures during instance running, the possible causes are as follows:

The maximum number of connections has been reached.

Memory or shards have been used up.

A high-availability (HA) switchover occurs, the database service becomes unavailable, a read-only replica switchover occurs, the read-only replica service becomes unavailable, or more.

For client errors, the possible causes are as follows:


Connection pool parameter setting is not proper.

Connections leak.

Troubleshooting

Step 1. Run `telnet` to Check Whether the Redis Port Can Be Accessed Normally

For most connection failure issues, run `telnet` in the command line tool to narrow down the cause of the error:



```
[root@VM-4-10-centos ~]# telnet 10.x.x.34 6379
Trying 10.x.x.34...
Connected to 10.x.x.34.
Escape character is '^]'.

```

As shown above, if there is successful connection prompt, the port of the Redis instance can be accessed normally. If there is an exception, go to [Step 2](#) to troubleshoot network issues.

Step 2. Check Whether It Is Caused by Network Configuration

To connect over the private network, the CVM and TencentDB for Redis must be under the same account and in the same VPC, or both in the classic network. The connection will fail in the following conditions:

If the CVM instance is in a VPC, while the Redis instance in the classic network, you are advised to switch the network type of the Redis instance from the classic network to VPC.

If the Redis instance is in a VPC, while the CVM instance in the classic network, you are advised to switch the network type of the CVM instance from classic network to VPC. For more information, see [Switching to VPC](#).

If the CVM and Redis are in different VPCs in the same region, you are advised to migrate the Redis instance to the VPC of the CVM instance.

If the CVM and Redis are in different VPCs in different regions, you are advised to create a [CCN](#) between the two VPCs.

If the CVM and Redis are in different VPCs under different accounts, you are advised to create a [CCN](#) between the two VPCs.

Step 3. Check Whether It Is a Security Group Issue

The CVM instance cannot connect to the Redis if their security group configuration is incorrect.

Incorrect CVM Security Group Configuration

To use the CVM to connect the Redis, you need to configure an outbound rule in the security group of the CVM instance. **If the target of the outbound rule is not "0.0.0.0/0" and the protocol port is not "ALL"**, the IP address and port of the Redis instance need to be added to the rule.

1. Go to the [Security Group](#) page in the CVM console and click the name of the CVM-bound security group to open its details page.
2. On the **Outbound rule** tab, click **Add Rule**.

Type: Select **Custom**.

Target: Enter the IP address or IP range of your Redis.

Protocol Port: Enter the private network port of the Redis.

Policy: Select **Allow**.

Incorrect Redis Security Group Configuration

To use the CVM to connect the Redis, you need to configure an inbound rule in the security group of the Redis instance. **If the source of the inbound rule is not "0.0.0.0/0" and the protocol port is not "ALL"**, the IP address and port of the CVM instance need to be added to the rule.

1. Go to the [Security Group](#) page in the CVM console and click the name of the Redis-bound security group to open its details page.
2. On the **Inbound rule** tab, click **Add Rule**.

Note that you also need to enable the IP address and port of the Redis in the inbound rule.

Type: Select **Custom**.

Source: Enter the IP address or IP range of your CVM.

Protocol Port: Enter the private network port of the Redis instance.

Policy: Select **Allow**.


Note:

The default private network port of Redis is 6379, and the port can be customized. If the default port is changed, the new port need be enabled in the inbound rule of the Redis security group.

If the default port 6379 of the Redis instance is used, it need to be enabled in the inbound rule of the Redis security group.


Step 4. Check Whether the Issue Is Caused by the Password

Run the `info` command. If the following information is displayed, the password of the Redis is correct.



```
[root@SNG-Qcloud /data/home/rickyu]# redis-cli -h 10.x.x.34 -p 6379 -a password
10.x.x.2:6379> info cpu
# CPU
used_cpu_sys:1623.176000
used_cpu_user:4649.572000
used_cpu_sys_children:0.000000
used_cpu_user_children:0.000000
```

If `NOAUTH Authentication required.` is displayed, the password is incorrect.

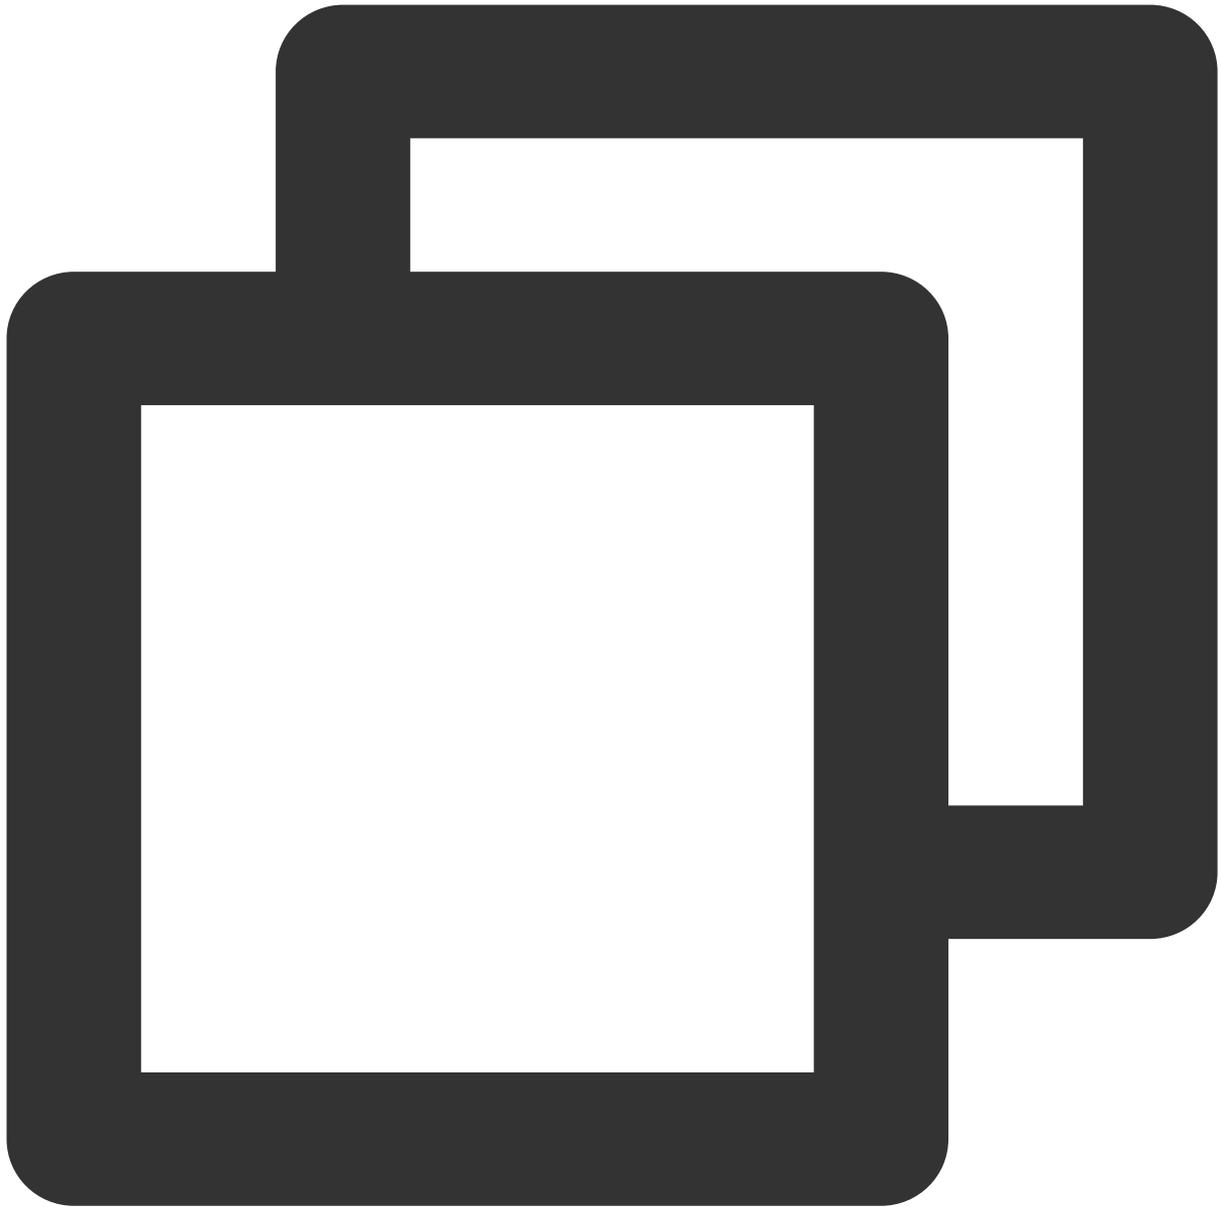


```
10.x.x.31:6379> info memory
NOAUTH Authentication required.
10.x.x.31:6379>
```

Log in to the [Redis console](#) and click an instance ID in the instance list to open the instance details page, where you can reset the password. For more information, see [Managing Account](#).


Step 5. Check Whether the Memory or Shards Are Full

If the following error message is displayed for the business:



```
"-READONLY You can't write against a read only slave.\\r\\n"
```

Log in to the [Redis console](#), find the target instance in the instance list, click the instance ID to enter the **System Monitoring** page, and select **Memory Utilization** in the **Metric** drop-down list to view the memory usage of the instance.



If memory is full, writes will fail. In this case, perform the following operations:

Expand the capacity immediately as instructed in [Changing Instance Specification](#).

Modify the database eviction policy as instructed in [Managing Instance Parameter](#). In **Parameter Settings**, set the

`maxmemory-policy` parameter to `allkeys-lru` or `volatile-lru`. For details, see [Setting Instance Parameters](#)

Note:

Instance data may be lost if the `allkeys-lru` eviction policy is adopted. Assess the impact before doing so.

Step 6. Check the Connection Quota

The **Connection Utilization** metric refers to the ratio of the number of TCP connections from the client to the instance to the maximum number of connections to the instance. If this metric is continuously high, the current database connection quota is insufficient, and the maximum number of connections needs to be adjusted.

Issue

The error message is as follows:

```
ERR max number of clients reached
```

Solution

1. Log in to the [Redis console](#). Above the instance list on the right, select the region. In the instance list, find the target instance. Click the instance ID in blue to enter the **Instance Details** page, click the **System Monitoring** tab, and then select the **Monitoring Metrics** tab to view the monitoring data. In the **View** drop-down list, select **Instance Monitoring** and **Connection Utilization** as the metric, and check whether it is continuously high in the monitoring view.
2. If the connection usage is continuously high, adjust the maximum number of connections and modify the specifications for the number of connections. For details, see [Adjusting the Number of Connections](#).

Step 7. Check Whether HA Switchover, Unavailable Database Service, Read-only Replica Switchover, or Unavailable Read-only Replica Service Occurs

If you find abnormal connections or a large number of access errors and slow queries at a certain point in time, and TCOP event alarms for those abnormal events are received, [Contact Us](#) for help.

For the configuration method of TCOP event alarms, see [Creating Event Rule](#).

Step 8. Confirm Whether the Configuration of the Jedis Connection Pool Is Correct If You Use It

Issue

If the number of available connections in the connection pool is used up and the old connections are not released in time, the newly created connection will fail, and the following error message will be displayed on the client.

```
JedisConnectionException: Could not get a resource from the pool
```

Solution

1. Use the following command on the client to confirm the number of connections currently accessing port 6379 of the instance. If this number is close to the `maxTotal` value configured in the connection pool, a connection failure will occur.

```
netstat -an | grep 6379 | grep ESTABLISHED | wc -l
```

2. See [Java Connection Sample](#), and check whether to call `jedis.close()` to release old connections to avoid connection leaks.

3. If all old connections have been released and the concurrent business volume is large, the `maxTotal` parameter value needs to be increased.

Note:

The `maxTotal` value of each client connection pool * The number of clients = The maximum number of connections for TencentDB for Redis

References

Viewing the Network Type and VPC Information

To enable connection between CVM and Redis over private network, they must be under the same account and in the same VPC, or both in the classic network.

Note:

If the **Network** fields in the instance lists both show **Classic Network** or **VPC**, the networks of the CVM and Redis are of the same type.

If the **Network** fields in the instance lists both show the same **VPC** (in the same region), it means that the CVM and TencentDB for Redis instances are in the same VPC.

Viewing the CVM Network Type

Log in to the [CVM console](#) and view the **Network** in the **Instance List**.

ID/Name	Monitoring	Status	Availability zone	Instance type	Instance configuration	Primary IPv4	Primary IPv6	Instance billing mode	Network
ins-...		Running	Guangzhou Zone 6	Compute C6	16-core 32GB 1Mbps System disk: CLOUD_PREMIUM Network: guangzhou			Monthly subscription Expire 5 days later	Prepaid

Viewing the Redis Network Type

Log in to the [Redis console](#) and view the **Network** in the **Instance List**.

Instance ID/Name	Monitoring/Status/Task	AZ	Network	Billing Mode	Architecture	Instance Edition	Used/Total	Creation Time
crs-...	Running	Guangzhou Zone 6		Pay as You Go	Redis 5.0Standard Architecture	Memory Edition	1GB/1GB(100.01%)	2024-03-09 10:16:27

Enabling Public Network Access

TencentDB for Redis now allows you to manually enable public network access in the console, to access Redis instance from the public network. For detailed directions, see [Configuring the Public Network Address](#).

Exception Analysis and Solution of Redisson Client Timeout Reconnection

Last updated : 2024-03-29 15:54:49

This article aims to briefly discuss the business blocking issues encountered by the Redisson client during network fluctuations and service failover, and proposes a simple and effective business layer solution to inspire solutions to similar problems.

Overview

When the Redisson client is used connect to a TencentDB for Redis instance, the VPC gateway fails and the failure lasts nearly 10 seconds. To restore the network connection, the faulty network device is isolated and the network link is recovered. Although the number of connections is restored to the normal business level, error reports for businesses affected by the network failure are still reported continuously. Only after the client is restarted and the system re-establishes the connection with the Redis database, business requests return to normal.

Exception Analysis

Based on the process of the exception, the scenario is replicated to access error logs and deeply analyze the internal working mechanism of Redisson connections, to locate potential issues.


Replicating the Issue

For the complex scenario of continuous network jitter, to facilitate testing and clear presentation, iptables rules are configured on the CVM to block access to Redis services, thereby simulating the effect of network jitter. Then, the block is lifted, the network link is restored, and error logs from business requests are collected.

1. To ensure that code clarity and ease of understanding, design a simple Redisson client example to connect to a Redis server, and execute a simple while loop upon successful connection to simulate a continuous task.

Note:

In the following code, set the Redis server's address, port, password, and timeout (3000ms) information in the `config-file.yaml` configuration file. Based on the configuration information, establish a connection to Redis. This example is based on Redisson 3.23.4, and different versions of Redisson may have some variations. You are advised to consult or inquire with Redisson's official for analysis.



```
package org.example;

// Import Redisson-related packages
import org.redisson.Redisson;
import org.redisson.api.RAtomicLong;
import org.redisson.api.RedissonClient;
import org.redisson.config.Config;
import org.redisson.connection.balancer.RoundRobinLoadBalancer;

// Import Java IO-related packages
import java.io.File;
```

```
import java.io.IOException;

public class Main {
    public static void main(String[] args) throws IOException {
        // Read Redisson's configuration information from a YAML file
        Config config = Config.fromYAML(new File("/data/home/test**/redissontest/sr

        RedissonClient redisson = Redisson.create(config);


        // Initialize a counter variable
        int i = 0;

        // Loop 1,000,000 times
        while(i++ < 1000000){
            // Access an atomic long object, where the key is the current counter's
            RAtomicLong atomicLong = redisson.getAtomicLong(Integer.toString(i));

            // Perform an atomic decrement by 1 on the value of the atomic long obj
            atomicLong.getAndDecrement();
        }

        // Close the Redisson client
        redisson.shutdown();
    }
}
```

2. After the client starts normally and runs for a short period, modify the iptables configuration to block connections to Redis.




```
sudo iptables -A INPUT -s 10.0.16.6 -p tcp --sport 6379 -m conntrack --ctstate NEW,
```

3. On the [Redis console](#), in the **System Monitoring** page, check the **Connections** metric. This metric shows a straight-line downward trend, as shown in the following figure.




4. Execute the following command to modify the iptables configuration, remove the connection block, and restore the network.



```
sudo iptables -D INPUT -s 10.0.16.6 -p tcp --sport 6379 -m conntrack --ctstate NEW,
```

5. After the network becomes normal, quickly perform reconnection on the client. On the [Redis console](#), in the **System Monitoring** page, peek at the metric again. The **Number of Connections** returns to the original level, but the overall requests are still abnormal, as shown in the following figure.



6. Obtain the abnormal log information, as shown in the following figure.

```

at java.base/java.lang.Thread.run(Thread.java:829)
[redisson-timer-4-1] ERROR org.redisson.client.handler.PingConnectionHandler - Unable to send PING command over channel
- R:10.0.16.6/10.0.16.6:6379]
org.redisson.client.RedisTimeoutException: Command execution timeout for command: (PING), params: [], Redis client: [ad
at org.redisson.client.RedisConnection.lambda$async$0(RedisConnection.java:256)
at io.netty.util.HashedWheelTimer$HashedWheelTimeout.run(HashedWheelTimer.java:715)
at io.netty.util.concurrent.ImmediateExecutor.execute(ImmediateExecutor.java:34)
at io.netty.util.HashedWheelTimer$HashedWheelTimeout.expire(HashedWheelTimer.java:703)
at io.netty.util.HashedWheelTimer$HashedWheelBucket.expireTimeouts(HashedWheelTimer.java:790)
at io.netty.util.HashedWheelTimer$Worker.run(HashedWheelTimer.java:503)
at io.netty.util.concurrent.FastThreadLocalRunnable.run(FastThreadLocalRunnable.java:30)
at java.base/java.lang.Thread.run(Thread.java:829)
[redisson-timer-4-1] ERROR org.redisson.client.handler.PingConnectionHandler - Unable to send PING command over channel
- R:10.0.16.6/10.0.16.6:6379]
org.redisson.client.RedisTimeoutException: Command execution timeout for command: (PING), params: [], Redis client: [ad
at org.redisson.client.RedisConnection.lambda$async$0(RedisConnection.java:256)
at io.netty.util.HashedWheelTimer$HashedWheelTimeout.run(HashedWheelTimer.java:715)
at io.netty.util.concurrent.ImmediateExecutor.execute(ImmediateExecutor.java:34)
at io.netty.util.HashedWheelTimer$HashedWheelTimeout.expire(HashedWheelTimer.java:703)
at io.netty.util.HashedWheelTimer$HashedWheelBucket.expireTimeouts(HashedWheelTimer.java:790)
at io.netty.util.HashedWheelTimer$Worker.run(HashedWheelTimer.java:503)
at io.netty.util.concurrent.FastThreadLocalRunnable.run(FastThreadLocalRunnable.java:30)
at java.base/java.lang.Thread.run(Thread.java:829)
[redisson-timer-4-1] ERROR org.redisson.client.handler.PingConnectionHandler - Unable to send PING command over channel
- R:10.0.16.6/10.0.16.6:6379]
org.redisson.client.RedisTimeoutException: Command execution timeout for command: (PING), params: [], Redis client: [ad
at org.redisson.client.RedisConnection.lambda$async$0(RedisConnection.java:256)
at io.netty.util.HashedWheelTimer$HashedWheelTimeout.run(HashedWheelTimer.java:715)
at io.netty.util.concurrent.ImmediateExecutor.execute(ImmediateExecutor.java:34)
at io.netty.util.HashedWheelTimer$HashedWheelTimeout.expire(HashedWheelTimer.java:703)
at io.netty.util.HashedWheelTimer$HashedWheelBucket.expireTimeouts(HashedWheelTimer.java:790)
at io.netty.util.HashedWheelTimer$Worker.run(HashedWheelTimer.java:503)
at io.netty.util.concurrent.FastThreadLocalRunnable.run(FastThreadLocalRunnable.java:30)
at java.base/java.lang.Thread.run(Thread.java:829)

```

Source Code Analysis


Based on the error messages, the code is analyzed one by one to locate the issue. Redisson does not attempt to re-access a new channel after encountering an exception during reconnection, leading to continuous request errors.

1. In the error message `Unable to send PING command over channel`, check the `PingConnectionHandler` code, as shown in the following figure. It shows that after the Ping reconnect detection fails with an error,

```
ctx.channel().close()
```

is executed to close the abnormal channel, and

`connection.getRedisClient().getConfig().getFailedNodeDetector().onPingFailed();` is executed to check the current connection client's status.




```
if(connection.getUsage() == 0 && future != null && (future.cancel(false) || cause(
    Throwable cause = cause(future);
    if (!(cause instanceof RedisRetryException)) {
        if (!future.isCancelled()) {
            log.error("Unable to send PING command over channel: {}", ctx.channel())
        }
        log.debug("channel: {} closed due to PING response timeout set in {} ms", c
            ctx.channel().close();
```

```
        connection.getRedisClient().getConfig().getFailedNodeDetector().onPingFailure
    } else {
        connection.getRedisClient().getConfig().getFailedNodeDetector().onPingSuccessful
        sendPing(ctx);
    }
} else {
    connection.getRedisClient().getConfig().getFailedNodeDetector().onPingSuccessful
    sendPing(ctx);
}
```

2. Further analyze the Connection source `RedisConnection connection =`

`RedisConnection.getFrom(ctx.channel());` and check the constructor for creating the connection

`RedisConnection()`, as shown in the following figure. This function updates the channel attributes and records the last use time of the connection but does not switch to a new channel for connection attempts.




```
public <C> RedisConnection(RedisClient redisClient, Channel channel, CompletableFu
    this.redisClient = redisClient;
    this.connectionPromise = connectionPromise;


    updateChannel(channel);
    lastUsageTime = System.nanoTime();

    LOG.debug("Connection created {}", redisClient);
}
// updateChannel updates the attributes of the Channel
```

```
public void updateChannel(Channel channel) {  
    if (channel == null) {  
        throw new NullPointerException();  
    }  
    this.channel = channel;  
    channel.attr(CONNECTION).set(this);  
}
```


3. If cluster information is updated, the following error message will be displayed.





```
ERROR org.redisson.cluster.ClusterConnectionManager - Can't update cluster stateo
```

4. In `ClusterConnectionManager.java`, you can locate the error message fragment related to updating the cluster. It shows that after the `Can't update cluster state` exception occurs, only the change in cluster status is checked, without retrieving a new connection from the connection pool for operation, and the operation of returning is directly performed.




```
private void checkClusterState(ClusterServersConfig cfg, Iterator<RedisURI> itera
    if (!iterator.hasNext()) {
        if (lastException.get() != null) {
            log.error("Can't update cluster state", lastException.get());
        }
        // Check for changes in cluster status
    }
    scheduleClusterChangeCheck(cfg);
    return;
}
.....
}
```

Solutions

Based on the above issue, adjust the business-side code. In case of an exceptional connection, actively interrupt the current connection thread, perform the Channel and connection configuration again, and attempt a reconnection.

Code Adjustment

Based on the code, when the issue was reproduced, optimize the following code to include operations for exceptional reconnection.



```
package org.example;
import org.redisson.Redisson;
import org.redisson.api.RAtomicLong;
import org.redisson.api.RedissonClient;
import org.redisson.config.Config;
import org.redisson.connection.balancer.RoundRobinLoadBalancer;
import java.io.File;
import java.io.IOException;


public class Main {
    public static void main(String[] args) throws IOException {
```

```
boolean connected = false;
RedissonClient redisson = null;
while (!connected) {
    try {
        Config config = Config.fromYAML(new File("/data/home/sharmaxia/redi
        redisson = Redisson.create(config);
        connected = true;
    } catch (Exception e) {
        e.printStackTrace();
        try {
            Thread.sleep(1000);
        } catch (InterruptedException ex) {
            Thread.currentThread().interrupt(); // Actively interrupt the c
        }
    }
}
int i = 0;
while (i++ < 1000000) {
    try {
        RAtomicLong atomicLong = redisson.getAtomicLong(Integer.toString(i)
        atomicLong.getAndDecrement();
    } catch (Exception e) {
        e.printStackTrace();
        try {
            Thread.sleep(1000);
        } catch (InterruptedException ex) {
            Thread.currentThread().interrupt(); // Actively interrupt the c
        }
        i--;
    }
}
redisson.shutdown();
}
```


Result Verification

Check the number of connections and the total number of requests and ensure that the business has returned to normal, to verify that adjustments to the code have resolved the issue of business disruptions caused by Redisson client timeout reconnection exceptions.

Number of Connections



Total Requests



Performance Troubleshooting and Fine-Tuning

Redis Instance Connection Failure

Last updated : 2023-06-06 15:38:52

Issue

Issue 1: Use a CVM instance to connect to a TencentDB for Redis instance at the private network address automatically assigned by the system as instructed in [Connecting to TencentDB for Redis Instance](#), but the connection fails.

Issue 2: Log in to the [TencentDB for Redis console](#), go to the instance list, click **Log In** in the **Operation** column of the target instance to redirect to DMC, connect to the instance, but the connection fails.

Possible Causes

For database connection failures for the first time, the possible causes are as follows:

Port error.

Network configuration error or incorrect security group configuration.

Password error.

For sudden connection failures during instance running, the possible causes are as follows:

The maximum number of connections has been reached.

Memory or shards have been used up.

A high-availability (HA) switch occurred, the database service became unavailable, a read-only replica switch occurred, or the read-only replica service became unavailable, etc.

For client errors, the possible causes are as follows:


Connection pool parameter setting is unreasonable.

Connections leaked.

Troubleshooting

Step 1. Run `telnet` to check whether the Redis port can be accessed normally

Run `telnet` in the command line tool to narrow down the cause of the error:



```
[root@VM-4-10-centos ~]# telnet 10.x.x.34 6379
Trying 10.x.x.34...
Connected to 10.x.x.34.
Escape character is '^]'.

```

As shown above, successful connection means that the port of the Redis instance can be accessed normally. If there is an exception, go to [Step 2](#) to troubleshoot network issues.

Step 2. Check whether it is caused by network configuration

To connect over the private network, the CVM and TencentDB for Redis instances must be under the same account and in the same VPC, or both in the classic network. The connection will fail in the following conditions:

If the CVM instance is in a VPC, while the Redis instance in the classic network, we recommend that you switch the network type of the Redis instance from classic network to VPC.

If the Redis instance is in a VPC, while the CVM instance in the classic network, we recommend that you switch the network type of the CVM instance from classic network to VPC. For more information, see [Switching to VPC](#).

If the CVM and TencentDB for Redis instances are in different VPCs in the same region, we recommend that you migrate the Redis instance to the VPC of the CVM instance.

If the CVM and TencentDB for Redis instances are in different VPCs in different regions, we recommend that you create a [CCN](#) between the two VPCs.

If the CVM and TencentDB for Redis instances are in different VPCs under different accounts, we recommend that you create a [CCN](#) between the two VPCs.

Step 3. Troubleshoot the security group

The CVM instance cannot connect to the TencentDB for Redis instance if their security groups are incorrect.

Incorrect CVM security group configuration

To use the CVM instance to access the Redis instance, you need to configure an outbound rule in the security group of the CVM instance. **If the target of the outbound rule isn't "0.0.0.0/0" and the protocol port isn't "ALL"**, the IP and port of the Redis instance should be added to the rule.

1. Go to the [Security Group](#) page in the CVM console and click the name of the CVM-bound security group to enter its details page.
2. On the **Outbound rule** tab, click **Add Rule**.

Type: Select **Custom**.

Target: Enter the IP or IP range of your Redis instance.

Protocol Port: Enter the private network port of the Redis instance.

Policy: Select **Allow**.

Incorrect Redis security group configuration

To use the CVM instance to access the Redis instance, you need to configure an inbound rule in the security group of the Redis instance. **If the source of the inbound rule isn't "0.0.0.0/0" and the protocol port isn't "ALL"**, the IP and port of the CVM instance should be added to the rule.

1. Go to the [Security Group](#) page in the CVM console and click the name of the Redis-bound security group to enter its details page.
2. On the **Inbound rule** tab, click **Add Rule**.

Note that you also need to open the IP and port of the Redis instance in the inbound rule.

Type: Select **Custom**.

Source: Enter the IP or IP range of your CVM instance.

Protocol Port: Enter the private network port of the Redis instance.

Policy: Select **Allow**.


Note

The Redis instance uses private network port 6379 by default and supports customizing its port. If the default port is changed, the new port should be opened in the inbound rule of the Redis security group.

If the default port 6379 of the Redis instance is used, it should be opened in the inbound rule of the Redis security group.


Step 4. Troubleshooting the password

Run the `info` command. If the following information is displayed, the password of the TencentDB for Redis instance is correct.



```
[root@SNG-Qcloud /data/home/rickyu]# redis-cli -h 10.x.x.34 -p 6379 -a password
10.x.x.2:6379> info cpu
# CPU
used_cpu_sys:1623.176000
used_cpu_user:4649.572000
used_cpu_sys_children:0.000000
used_cpu_user_children:0.000000
```

If `NOAUTH Authentication required.` is displayed, the password is incorrect.




```
10.x.x.31:6379> info memory  
NOAUTH Authentication required.  
10.x.x.31:6379>
```

Log in to the [TencentDB for Redis console](#) and click an instance ID in the instance list to enter the instance details page, where you can reset the password. For more information, see [Managing Account](#).

Step 5. Check whether the memory or shards are used up

If the business reports the following error message:



```
"-READONLY You can't write against a read only slave.\\r\\n"
```

Log in to the [TencentDB for Redis console](#), find the target instance in the instance list, click the instance ID to enter the **System Monitoring** page, and select **Memory Utilization** in the **Metric** drop-down list to view the memory utilization of the instance.

If memory is used up, writes will fail. In this case, perform the following operations:
Expand the capacity immediately as instructed in [Changing Instance Specification](#).

Modify the database eviction policy as instructed in [Managing Instance Parameter](#). In **Parameter Settings**, set the `maxmemory-policy` parameter to `allkeys-lru` or `volatile-lru`.

Note

Instance data may be lost if the `allkeys-lru` eviction policy is adopted. Assess the impact before doing so.

Step 6. Check the connection quota

The **Connection Utilization** metric refers to the ratio of the number of TCP connections from the client to the instance to the maximum number of connections to the instance. If this metric stays high, the current database connection quota is insufficient, and the maximum number of connections needs to be adjusted.

Issue

The error message is as follows:

```
ERR max number of clients reached
```

Solution

1. Log in to the [TencentDB for Redis console](#), Above the instance list on the right, select the region. In the instance list, find the target instance. Click the instance ID in blue to enter the **Instance Details** page, click the **System Monitoring** tab, and then select the **Monitoring Metrics** tab to view the monitoring data. In the **View** drop-down list, select **Instance Monitoring** and **Connection Utilization** as the metric, and check whether it stays high in the monitoring view.
2. If the connection utilization stays high, adjust the maximum number of connections and modify the specifications for the number of connections. For details, see [Adjusting the Number of Connections](#).

Step 7. Check whether any of the following occurred: HA switch, unavailable database service, read-only replica switch, or unavailable read-only replica service

If you find abnormal connections or a large number of access errors and slow queries at a certain point in time, and you also receive TCOP event alarms for those abnormal events, [contact us](#) for help.

For the configuration method of TCOP event alarms, see [Creating Event Rule](#).

Step 8. Confirm whether the configuration of the Jedis connection pool is correct if you use it

Issue

If the number of available connections in the connection pool is used up and the old connections are not released in time, the newly created connection will fail, and the client will prompt the following error message.

```
JedisConnectionException: Could not get a resource from the pool
```

Solution

1. Use the following command on the client to confirm the number of connections currently accessing port 6379 of the instance. If this number is close to the `maxTotal` value configured in the connection pool, a connection failure will occur.

```
netstat -an | grep 6379 | grep ESTABLISHED | wc -l
```

2. See [Java Connection Sample](#), and check whether to call `jedis.close()` to release old connections to avoid connection leaks.

3. If all old connections have been released and the concurrent business volume is large, the `maxTotal` parameter value needs to be increased.

Note

The `maxTotal` value of each client connection pool * the number of clients = the maximum number of connections for TencentDB for Redis.

References

Viewing the network type and VPC information

To enable connection between CVM and TencentDB for Redis instances over private network, they must be under the same account and in the same VPC, or both in the classic network.

Note

If the **Network** fields in the instance lists both show **Classic Network** or **VPC**, it means that the networks of the CVM and TencentDB for Redis instances are of the same type.

If the **Network** fields in the instance lists both show the same **VPC** (in the same region), it means that the CVM and TencentDB for Redis instances are in the same VPC.

Viewing the CVM network type

Log in to the [CVM console](#) and view the **network information** in the instance list.

Viewing the Redis network type

Log in to the [TencentDB for Redis console](#) and view the **network information** in the instance list.

Enabling public network access

TencentDB for Redis now allows you to manually enable public network access in the console, so that instances can be accessed over public network. For detailed directions, see [Configuring the Public Network Address](#).

High CPU Utilization

Last updated : 2023-05-23 10:30:16

Increasing CPU utilization will affect the throughput of the entire instance cluster, and may even cause application blocking and timeout interruptions. When the average CPU utilization exceeds 60% or the average peak CPU utilization stays higher than 90% for more than 5 minutes, you need to check the specific cause in time and solve the problem quickly to ensure business stability and availability.

Symptom

Symptom 1: You received an alarm about high CPU utilization.

Symptom 2: The value of the CPU utilization metric was high.

Symptom 3: The overall throughput became lower and the response slower.

Troubleshooting

No.	Possible Causes	Cause Analysis	Troubleshooting Methods	Solution
1	Non-persistent connections are established frequently.	A large number of resources of the instance are consumed in processing frequent non-persistent connections, resulting in high CPU utilization and a lot of connections. However, the QPS (cluster accesses per second) did not meet expectations.	With DBbrain's Performance Optimization feature, you can analyze the real-time session statistics view and data of the database instance to check whether there is a sudden increase in the number of connections. For directions, see Real-Time Session .	Change the non-persistent connection to a persistent connection, for example, for example, by using the JedisPool connection pool. For specific code samples, see Jedis Client .
2	There are commands	As Redis	Executing high-	Avoid using complex

	<p>with high time complexity, such as <code>sort</code> , <code>sunion</code> , and <code>zunionstore</code> .</p>	<p>executes commands in a single thread, the execution of complex commands may block other commands. The higher the time complexity of the command, the more resources will be consumed during execution and the CPU utilization will increase.</p>	<p>complexity commands is usually time-consuming and will generate slow logs accordingly. You can use DBbrain's Performance Optimization feature to perform Slow Log Analysis, and check whether there are high-complexity commands in the slow log information list. For directions, see Slow Log Analysis.</p>	<p>commands to get large amounts of data. Try to just use them on a limited amount of data instead, so Redis can process it quickly and return the results.</p>
<p>3</p>	<p>The read and write load is too high due to big keys or a large number of accesses to hot keys.</p>	<p>Hot keys refer to those key values that are accessed very frequently over a period of time, such as trending news, popular live streaming, and flash sales. In these cases, a large amount of access traffic is concentrated in one instance, reaching the processing capacity of a single instance and causing an increase in CPU utilization.</p>	<p>With DBbrain's Performance Optimization feature, you can perform Hot Key Analysis to quickly find hot keys with high access frequency. For directions, see Hot Key Analysis.</p>	<p>You can split hot keys of complex data structures into several new keys and distribute them across Redis nodes to reduce the pressure. For example, if a two-level hash hot key has a lot of hash elements, you can split it.</p>
<p>4</p>		<p>A big key is one that has big value and takes up a large space. When read or</p>	<p>With DBbrain's Performance Optimization feature, you can perform Big Key Analysis and monitor</p>	<p>If the value of a big key is too large, you can split the big key into multiple key-values so that multiple Redis nodes will</p>

		<p>deletion operations related to big keys are performed, bandwidth and CPU will be seriously occupied.</p>	<p>and analyze the memory usage of big keys in the database. For directions, see Memory Analysis.</p>	<p>share the pressure. If there are too many keys, you can store them in a hash structure.</p>
<p>5</p>		<p>The read load is too high and the resource limit has been reached.</p>	<p>With DBbrain's Performance Optimization feature, you can check the Performance Trends view of the database instance, analyze QPS and read/write request metrics, and check whether the high CPU utilization is caused by excessive read load or write load. For directions, see Performance Trends.</p>	<p>If the read load is too high, you can share the read load by increasing the number of replicas as instructed in Changing Instance Specification. You need to enable read-only replicas and transfer the read requests of the current instance to the read-only replica node, so as to realize the elastic expansion of the read capability and improve the read and write performance. For directions, see Enabling/Disabling Read/Write Separation.</p>
<p>6</p>		<p>The write load is too large, and the memory specification is insufficient.</p>		<p>If the write load is too high, you can share the write load by increasing the number of shards as instructed in Changing Instance Specification. If the instance is on a standard architecture, you need to upgrade the standard architecture to a cluster architecture to improve CPU processing capabilities. For directions, see Upgrading Instance Architecture. You need</p>

				to check compatibility before upgrading to the cluster architecture. For more information, see Check on Migration from Standard Architecture to Cluster Architecture .
7	A large number of keys expire at the same time.	The key expiration time is set at the same time point. When the expiration time point is reached, Redis will take up a lot of CPU resources to process the elimination thread of those expired keys, causing a momentary lagging.	On the System Monitoring page in the console, you can check the Expired Keys metric to confirm whether a large number of keys expire at a certain point in time. For directions, see Update Notes of Monitoring at Five-Second Granularity .	You can disperse the expiration time of those keys in the business logic to avoid them all expiring at the same time.
8	Databases are frequently switched by executing the <code>select</code> command.	Frequent database switch leads to excessive resource overhead.	With DBbrain's Latency Analysis feature, you can check the monitoring data of the <code>select</code> command in Command Word Analysis , and confirm whether there are many <code>select</code> requests. For directions, see Command Word Analysis .	For the storage of different businesses, if some of them require frequent database switch, we recommend that you store them separately. For the storage of the same businesses, if the key names do not conflict, you can consider storing them in the same database to reduce the number of <code>select</code> request operations.

High Outbound Traffic

Last updated : 2023-05-23 10:30:54

Symptom

Symptom 1: You received an alarm about the outbound traffic being restricted because the outbound traffic metric reached the maximum value allowed.

Symptom 2: The response latency increased.

Possible Causes

Big keys.

If the value of the request key is big, it is easy to cause the problem of high outbound traffic.

Pipeline requests.

Pipeline technology combines multiple requests into one request and send it to the database server for processing, receives all command execution results, and then returns to the upper-layer business. If there are many requests at a time, the outbound traffic will be too high. For more information, see [Redis pipelining documentation](#).

Batch query, such as mget, hmget or hgetall, etc.

The number of keys that are queried in a single batch is large, causing the outbound traffic to be too high.

The instance configuration specification is insufficient.

As business and data volume rise, the instance's outbound traffic has reached its upper limit and can't meet the current business traffic demands.

Solutions

Step 1. Troubleshoot big key issues

1. Log in to the [TencentDB for Redis console](#), use DBbrain's **Performance Optimization** feature to check the big key, and create a big key analysis task, and optimize the big key based on the analysis report. For details, see [Memory Analysis](#).
2. If there is a big key, you can reduce the access to the big key without affecting your business. If the value is too large, you can split the object into multiple key-values so that multiple Redis instances will share the pressure. If there are too many keys, you can store multiple keys in a hash structure.

Note

In order to prevent the generation of big keys, we recommend that you refer to the following suggestions when designing value.

The string big keys should be under 10 KB in size, and the number of hash, list, set, and zset elements should not exceed 5000.

For non-string big keys, it is recommended to use `hscan`, `sscan`, and `zscan` to delete them progressively.

Step 2. Check whether the business uses pipeline

To check the business code logic, it is recommended not to use **pipeline** or reduce the number of command requests in each **pipeline**. If you are unsure whether your business has used the pipeline, [submit a ticket](#) for troubleshooting.

Step 3. Check the number of keys in a single query

Log in to the [TencentDB for Redis console](#), use DBbrain's **Performance Optimization** feature to analyze the change trend of the performance metrics **Key Requests** and **Mget Executions**, and check whether there is a sudden increase in key requests. For directions, see [Performance Trends](#).

Generally, it is recommended that the number of keys or elements in a single operation should not exceed 50. If the value is relatively large, you need to reduce it.

Step 4. Upgrade the instance specification

1. Increase the number of replicas as instructed in [Changing Instance Specification](#), enable read-only replicas to share the read load, and transfer the read requests of the current instance to the read-only replica nodes to achieve elastic expansion of read capabilities and improve network traffic performance. For specific operations, see [Enabling/Disabling Read/Write Separation](#).

2. Increase the number of shards as instructed in [Changing Instance Specification](#), and the system will allocate more standard bandwidth. If the instance is on a standard architecture, you need to upgrade the standard architecture to the cluster architecture to improve CPU processing capabilities. For directions, see [Upgrading Instance Architecture](#). It is necessary to check the compatibility before upgrading to the cluster architecture. For more information, see [Check on Migration from Standard Architecture to Cluster Architecture](#).

Step 5. Adjust the bandwidth

If the outbound traffic is still high after upgrading the instance specification, you can adjust the bandwidth to its maximum value. Increasing bandwidth is currently free of charge. For directions, see [Bandwidth Adjustment](#).

High Memory Utilization

Last updated : 2023-05-23 10:49:31

Symptom

Symptom 1: You received an alarm about too high memory utilization.

Symptom 2: On the **System Monitoring** page in the [TencentDB for Redis console](#), you can view that there are sudden rises in **Memory Utilization**, **Key Evictions**, and **P99 Response Latency**.

Symptom 3: The error message `command not allowed when used memory > 'maxmemory'` was prompted when the program writes data.

Redis Memory Usage

The data that usually occupies memory in TencentDB for Redis includes the following:

Object memory: User data area, that is, the actual stored value information.

Buffer memory: Including client input and output buffers, and master-replica sync replication buffers.

Note

When performing Range operations or big keys on the client side, the memory occupied by input buff and output buff will increase, affecting the buffer and even causing OOM.

Memory fragmentation: A large number of update operations, such as append and setrange; a large number of expired keys are deleted, and the released space cannot be effectively used.

Link memory: The memory consumed for creating subprocesses, which is generally relatively small.

Troubleshooting the Issue

No.	Possible Cause	Troubleshooting Method	Solution
1	The increase in the volume of written data leads to an increase in memory usage. The key is not set in the TTL policy.	1. Log in to the TencentDB for Redis console , click an instance ID, and enter the instance details page. 2. Select the System Monitoring tab to view the corresponding	If the memory usage is proportional to the total number of keys, the possible causes are as follows. The memory may be increased due to the normal data written by the business. Evaluate your business needs and expand the Redis instance in time. For directions, see Changing Instance Specification .

		<p>monitoring metrics of the TencentDB for Redis instance, including Memory Utilization, Total Keys, Expired Keys, and Evicted Keys. Then, you can analyze whether the memory utilization is consistent with the fluctuation trend of the number of keys.</p>	<p>The memory may increase due to the accumulation of keys due to the failure to set a reasonable eviction policy and deletion frequency for expired keys. The eviction policy setting is unreasonable, which may easily cause invalid keys to occupy too much space. To modify the parameters related to key expiration time, eviction policy, and deletion frequency, reconfigure the <code>maxmemory-policy</code> and <code>hz</code> parameters on the Parameter Settings page in the console. For directions, see Setting Instance Parameters. If the deletion frequency for expired keys is set too high, that is, if the value of <code>hz</code> increases, more CPU resources will be occupied. The <code>hz</code> value should not be too large. You can adjust it as needed.</p>
<p>2</p>	<p>Input buffer overflow is caused by writing a big key caused the . Input buffer overflow is caused by reading a big key, requesting a large number of commands to return results, or executing a monitor command.</p>	<p>Analyze memory skew issues with the Performance Optimization feature of DBbrain. For more information, see Memory Analysis.</p>	<p>Perform value splitting and optimization for abnormally large keys. For more information, see Hot Key and Big Key.</p>

High Number of Total Requests

Last updated : 2023-05-23 10:33:40

Symptom

Symptom 1: the QPS value was high.

Symptom 2: the response latency increased.

Symptom 3: connection timeout occurred.

Possible Causes

The business needs to be optimized.

The instance configuration needs to be upgraded.

Solutions

Check the node load: for the cluster architecture, check the node load. If the QPS of only one or a few nodes exceeds the alarm threshold, there may be a hot key; if the QPS of most nodes is high, the overall load of the TencentDB for Redis instance is high, in which case the instance configuration needs to be upgraded.

Check slow queries: you can check whether there are slow queries in the console, and if so and the slow queries match the time when the problem occurred, the problem may be caused by a big key.

Check the CPU utilization: you can check whether the CPU utilization is too high, and if so, the machine resources may be insufficient, in which case the instance configuration needs to be upgraded.

If your business requires optimization, you can optimize it in terms of hot keys and big keys. If the instance configuration requires upgrade, you can enable read/write separation and add more shards to meet your current business needs.

Troubleshooting the Issue

Optimizing your business

1. Log in to the [TencentDB for Redis console](#) and click an instance ID in the instance list to enter the instance management page.
2. On the **System Monitoring** tab, check whether QPS is high or whether there are unexpected hot keys.

3. After troubleshooting abnormal access, optimize your business logic:

Hot keys: You can split hot keys of complex data structures into several new keys and distribute them across Redis nodes to reduce the pressure. For example, if a two-level hash hot key has a lot of hash elements, you can split it.

Big keys: If the value is too large, you can split the object into multiple key-values so that multiple Redis nodes will share the pressure. If there are too many keys, you can store multiple keys in a hash structure.

Instance upgrade


Heavy read load

[Add replicas](#) and enable [read/write separation](#) to share the read load.

Note

Confirm that your business allows inconsistent data before enabling read/write separation, because after it is enabled, inconsistent data may be read from the replica node and the master node (the replica node lags behind the master node). For more information, see [Changing Instance Specification](#).

1. Log in to the [TencentDB for Redis console](#), find the target instance in the instance list, and select **Configure > Add Replica** in the **Operation** column.



Instance ID/Name	Monitoring/Status/Task	Project	AZ	Network	Billing Mode	Architecture	Instance Edition	User
[Redacted]	Running	Default Project	Guangzhou Zone 3	Default-VPC - Default-Subnet 10.0.0.13:6379(Ipv4)	Pay as You Go --	Redis 6.2Standard Architecture	Memory Edition	34.9

2. In the pop-up window, adjust the configuration and click **OK**.

3. Return to the instance list. After the instance becomes **Running** in the status, you can use it normally.

Heavy write load

Cluster Architecture

Log in to the [TencentDB for Redis console](#), find the target instance in the instance list, and select **Configure > Add Shard** in the **Operation** column.

Note

After the configuration is adjusted, the instance will be charged at the price of the new configuration.

When shards are added, the system will automatically balance the slot configuration and migrate data, which may fail in rare cases. We recommend that you perform such operations during off-peak hours to avoid the impact of migration on business access.

Add shards as needed: Each shard supports a QPS of 80,000 to 100,000.

Instance ID/Name	Monitoring/Status/Task	Project	AZ	Network	Billing Mode	Architecture	Instance Edition	Us
	  Running	Default Project	Guangzhou Zone 3	Default-VPC - Default-Subnet 10.0.0.13:6379(IPv4)	Pay as You Go --	Redis 6.2Cluster Architecture	Memory Edition	34

Standard Architecture

Upgrade the instance from standard architecture to cluster architecture to improve the processing power of CPU.

Before the upgrade, you need to check the compatibility. For more information, see [Check on Migration from Standard Architecture to Cluster Architecture](#).

1. Log in to the [TencentDB for Redis console](#) and click an instance ID in the instance list to enter the instance details page.
2. In the **Specs Info** section, click **Upgrade Architecture**.
3. After the upgrade is completed, go to the instance list and select **Configure > Add Shard** in the **Operation** column.

Note

If the problem persists, [contact us](#) for assistance.

Execution Error

Last updated : 2023-05-23 10:45:03

Symptom

Error 1: an execution error occurred.

Error 2: the effect is not as expected.

Possible Causes

The command is entered incorrectly.

If the command is entered correctly, the cause may be that:

There is a deviation in business logic understanding.

The memory is full.

Solutions

If the command is entered incorrectly, enter the correct command.

If the command is entered correctly, locate the problem by referring to [Error Codes](#). Common problems include the following:

Lua script execution failed.

The cluster command is executed without the correct node ID.

The memory is full.

Troubleshooting the Issue

For errors caused by the use of commands, you can fix them according to the corresponding descriptions and business logic.

For execution errors caused by full memory, you can fix them as instructed in [High Memory Utilization](#).

Note

If the problem persists, [contact us](#) for assistance.