Tencent Cloud Distributed Cache Release Notes and Announcements

©2013–2026 Tencent Cloud. All rights reserved.

Copyright Notice

©2013-2026 Tencent Cloud. All rights reserved.

The complete copyright of this document, including all text, data, images, and other content, is solely and exclusively owned by Tencent Cloud Computing (Beijing) Co., Ltd. ("Tencent Cloud"); Without prior explicit written permission from Tencent Cloud, no entity shall reproduce, modify, use, plagiarize, or disseminate the entire or partial content of this document in any form. Such actions constitute an infringement of Tencent Cloud's copyright, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

Trademark Notice

Tencent Cloud

This trademark and its related service trademarks are owned by Tencent Cloud Computing (Beijing) Co., Ltd. and its affiliated companies("Tencent Cloud"). The trademarks of third parties mentioned in this document are the property of their respective owners under the applicable laws. Without the written permission of Tencent Cloud and the relevant trademark rights owners, no entity shall use, reproduce, modify, disseminate, or copy the trademarks as mentioned above in any way. Any such actions will constitute an infringement of Tencent Cloud's and the relevant owners' trademark rights, and Tencent Cloud will take legal measures to pursue liability under the applicable laws.

Service Notice

This document provides an overview of the as-is details of Tencent Cloud's products and services in their entirety or part. The descriptions of certain products and services may be subject to adjustments from time to time.

The commercial contract concluded by you and Tencent Cloud will provide the specific types of Tencent Cloud products and services you purchase and the service standards. Unless otherwise agreed upon by both parties, Tencent Cloud does not make any explicit or implied commitments or warranties regarding the content of this document.

Contact Us

We are committed to providing personalized pre-sales consultation and technical after-sale support. Don't hesitate to contact us at 4009100100 or 95716 for any inquiries or concerns.

Contents

Release Notes and Announcements
- Release Notes
- Announcements
  - Announcement of Full Decommissioning of TencentCloud API 2.0
  - Classic Network Will Be Disused
  - API for Querying Instance Big Key Will Be Disused
  - Announcement of TencentDB for Redis 1-Minute Granularity Discontinuation

©2013-2026 Tencent Cloud. All rights reserved.

Release Notes and Announcements
Release Notes

Last updated: 2024-02-26 14:28:02

June 2023

Update	Description	Document
Commands added	The command compatibility of version 6.2 are enhanced, commands such as LPOP/RPOP COUNT are now supported, which provides smooth migration capabilities for businesses.	Command Compatibility Across Versions
Forced promotion supported in global replication	Global replication supports immediate switching of instance roles within a replication group, catering to various scenarios prioritizing business availability and reliability.	Switching instance role

May 2023

Update	Description	Document
Enhanced DTS capabilities	Support for migration from AWS Database to TencentDB for Redis.	Migration with DTS

February 2023

Update	Description	Document
Configure backup file download rules	To safeguard against unauthorized access to database backup data, TencentDB for Redis enables users to designate specific servers for downloading backup files within the permitted private network.	Configure Backup File Download Rules
Selection of billing methods	When cloning a complete new instance based on the current backup file, you can flexibly choose between a	Cloning Instances

©2013-2026 Tencent Cloud. All rights reserved.

for clone instances allowed	prepaid monthly subscription or a pay-as-you-go billing mode.	

November 2022

Update	Description	Document
Redis 6.2 released	Redis 6.2 is fully released, supporting both standard and cluster architectures, read/write separation, seamless scaling, and offering ultra-high availability.	Product Series
Supported exporting parameters as a file	On the **Parameter Settings** page in the console, you can export the parameter file of the current instance to your local PC.	Manage Instance Parameters

October 2022

Update	Description	Document
Billing mode conversion supported	Supports conversion from pay-as-you-go to monthly subscription billing.	Switch from Pay-as-you-go to Monthly Subscription
Optimized the instance list in the console	• The **Deployment Mode** column is added to the instance list to indicate whether an instance is deployed across AZs. • When exporting the instance list information, you can aggregate tags for display, which are separated by |.	View Instance Information

September 2022

Update	Description	Document
Added monitoring metrics	The **Max Connections Utilization of Node** metric is added to monitor the maximum connection utilization of all proxy nodes in an instance.	Monitoring at Five-Second Granularity

August 2022

Update	Description	Document

©2013-2026 Tencent Cloud. All rights reserved.

Supported querying slow logs of read-only replicas	TencentDB for Redis supports slow query of read-only replica to help you troubleshoot performance problems of read-only operations.	-
Supported smooth scale-in to the 256 MB specification	TencentDB for Redis in standard architecture supports smooth reduction of memory capacity from 1 GB or higher to 256 MB.	Changing Instance Specification

July 2022

Update	Description	Documentation
Optimized statistical algorithm for traffic and throttling that triggers monitoring metrics	Optimized inbound/outbound traffic and the statistical algorithm for inbound traffic throttling and outbound traffic throttling triggers.	-
Supported the 256 MB specification for the standard architecture	TencentDB for Redis 4.0 and 5.0 Memory Edition (Standard Architecture) instances now offer the minimum memory specification of 256 MB in more Availability Zones.	Performance

June 2022

Update	Description	Documentation

©2013-2026 Tencent Cloud. All rights reserved.

Supported SSL encryption	TencentDB for Redis supports SSL encryption to implement encrypted data transfer.	SSL Encryption

May 2022

Update	Description	Documentation
Supported minor version upgrade	TencentDB for Redis supports minor version upgrades, assisting users in updating their instances to newer versions, thereby experiencing a wealth of enhanced features.	Upgrading Instance Version
Supported proxy upgrade	TencentDB for Redis periodically releases minor versions of the proxy to enhance database features or fix known defects.	Upgrading Proxy

April 2022

Update	Description	Documentation
Upgraded the global replication feature	Multi-AZ instances can be added to a global replication group, and existing instances can be quickly upgraded to global replication instances.	Creating Global Replication Group

February 2022

Update	Description	Documentation
Supported architecture upgrade for multi-AZ deployed instances	Multi-AZ deployed instances can be upgraded from the standard architecture to cluster architecture.	Upgrading Instance Architecture

©2013-2026 Tencent Cloud. All rights reserved.

Supported instance clone for multi-AZ deployed instances	Multi-AZ deployed instances allow you to clone a complete instance from a backup file.	Cloning Data

January 2022

Update	Description	Documentation
Global Replication	Global replication enables the addition and removal of instances, switching of instance roles, and deletion of replication groups.	Global Replication Overview
Supported custom backup retention period	You can apply for customizing the retention period of backup files through an allowlist.	Data Backup
Added the API for failure simulation	Instances within the same availability zone now support the invocation of the API interface KillMasterGroup for fault simulation testing.	KillMasterGroup
Added the API for master-replica switch	Instances within the same availability zone support invoking the ChangeReplicaToMaster API interface to elevate a replica (group) to a primary role, achieving the purpose of primary-secondary switching.	ChangeReplicaToMaster

December 2021

Update	Description	Documentation
Supported the 256 MB memory specification for	The minimum memory specification of TencentDB for Redis 4.0 and 5.0 Memory Edition (Standard Architecture) instances is 256 MB. Currently, this is supported only in Shanghai Zone 5, Beijing Zone 6, and Guangzhou Zone 6.	Performance

©2013-2026 Tencent Cloud. All rights reserved.

Tencent DB for Redis 4.0 and 5.0 Memory Edition (Standard Architecture)		

November 2021

Update	Description	Documentation
Supported system default templates	TencentDB for Redis now provides default parameter templates applicable to Redis 2.8 standard architecture, Redis 4.0 standard architecture, Redis 4.0 cluster architecture, Redis 5.0 standard architecture, and Redis 5.0 cluster architecture. You can view them on the parameter template page in the console.	Applying Parameter Templates
Supported enabling public network access	TencentDB for Redis now supports enabling public network access, allowing Cloud Virtual Machines and cloud databases in different VPCs to directly enable public network addresses in the Redis console for cross-public network connections.	Configuring Public Network Address

July 2021

Update	Description	Documentation
Supported parameter template	In addition to the system parameter templates provided by TencentDB for Redis, you can create custom parameter templates to configure parameters in batches according to your specific business scenarios.	Applying Parameter Templates

June 2021

Update	Description	Documenta

©2013–2026 Tencent Cloud. All rights reserved.

		tion
Supported automatic failover	TencentDB for Redis supports the automatic failover feature for proxy nodes and Redis server nodes to ensure service availability.	Failover
Auto-failback supported	TencentDB for Redis provides the automatic failback feature for instances deployed across AZs. After the feature is enabled, if the master node is switched from the master AZ or master node group (cluster architecture) to another AZ or group after a failover occurs, it will be automatically switched back, simplifying subsequent Ops operations.	Auto-Failback
Manually promote the master node	TencentDB for Redis offers a manual promotion feature for instances deployed across multiple AZs, supporting both standard and cluster architectures. You can deploy the master node to a specified AZ or replica (group) based on your business requirements.	Manually Promote Master Node
Nearby Access	In a multi-AZ deployment scenario, TencentDB for Redis offers the ability to access services with reduced latency by providing proximity-based access.	Read local nodes only

March 2021

Update	Description	Documentation
Redis monitoring upgraded to five-second granularity across the network	TencentDB for Redis monitoring is being progressively upgraded to a five-second granularity across the network. This upgrade will involve changes to cloud monitoring data access, monitoring metrics, Proxy, and alarm configurations.	Monitoring Upgrade and Alarm Migration Announcement

12/2020

Update	Description	Documentation

©2013-2026 Tencent Cloud. All rights reserved.

Redis Hybrid Storage Edition has been renamed to TencentDB for Tendis.	Redis Hybrid Storage Edition has been officially renamed to TencentDB for Tendis, featuring a dedicated product entry and an independent console.	TencentDB for Tendis
Supports Multi-AZ deployment	TencentDB for Redis supports deploying replicas across multiple availability zones within the same region, providing higher availability and disaster recovery capabilities compared to single availability zone instances (where the master and replica nodes are in the same availability zone).	Multi-AZ

September 2020

Update	Description	Documentation
Supports five-second monitoring granularity	The monitoring granularity of Redis has been upgraded from one minute to five seconds, providing a comprehensive enhancement to Redis monitoring in the new version.	Monitoring at Five-Second Granularity

07/2020

Update	Description	Documentation
Supported version upgrade	TencentDB for Redis now supports version upgrade, so you can upgrade from a lower version of Standard Edition to a higher one, including from 2.8 to 4.0, from 2.8 to 5.0, and from 4.0 to 5.0.	Upgrading Instance Version

©2013-2026 Tencent Cloud. All rights reserved.

Supported architecture upgrade	TencentDB for Redis now supports quick upgrade from standard architecture to cluster architecture to help your business expand the performance and capacity with speed.	Upgrading Instance Architecture

June 2020

Update	Description	Documentation
Launched TencentDB for Redis Hybrid Storage Edition	TencentDB for Redis Hybrid Storage Edition is officially launched, offering 100% compatibility with Redis protocols, reducing memory costs by up to 80%, and providing hot data performance consistent with the Memory Edition. It achieves a perfect balance between compatibility, performance, and cost.	Hybrid Storage Edition (Cluster Architecture)

April 2020

Update	Description	Documentation
Supported Redis 5.0 in DTS	Migrate data and upgrade instances to Redis 5.0 using Data Transmission Service (DTS).	Migration with DTS Version Upgrade with DTS

March 2020

Update	Description	Documentation
Supported the service unavailable alarm feature	TencentDB for Redis now supports the service unavailable alarm feature. In addition, users can also set alarms for instance master/replica switch, service unavailability, read-only replica failover, and read-only replica unavailability.	Configure Alarm
Supported adjusting bandwidth	You can now adjust instance network bandwidth on the instance details page in the Redis console. The bandwidth specifications for Cluster Edition have been	-

©2013-2026 Tencent Cloud. All rights reserved.

in the console	fully upgraded, with shard specifications upgraded to 384 Mbps.	
Updated the monitoring view for Redis 2.8	To provide more accurate monitoring information, we have updated the monitoring information of version 2.8 instances to a cluster view. If you are using APIs to obtain monitoring data from Cloud Monitor, you need to change the view in your code from redisuuid to cluster.	–

February 2020

Update	Description	Documentation
Upgraded the default number of databases in an instance to 256	For Redis 2.8, 4.0, and 5.0 Standard and Cluster Editions, the default number of databases in newly purchased instances has been upgraded from 16 to 256.	Memory Edition (Standard Architecture) Memory Edition (Cluster Architecture)

01/2020

Update	Description	Documentation
Released Redis 5.0	Redis 5.0 Standard and Cluster Editions have been released, supporting the latest STREAM data structure, new ZSET commands ZPOPMIN and ZPOPMAX, retaining all features of Redis 4.0, and offering continuous scaling, 4 TB of storage capacity, and tens of millions of concurrent QPS capabilities.	Memory Edition (Standard Architecture) Memory Edition (Cluster Architecture)

October 2019

Update	Description	Documentation

©2013–2026 Tencent Cloud. All rights reserved.

Supports Cloud Access Management	Support creating policies via Cloud Access Management, allowing sub-accounts to access the resources and permissions they require.	Cloud Access Management
Supported account management in the console	TencentDB for Redis provides read-write permission control and routing policy control through the account mechanism, which helps meet the needs of business permission management in complex scenarios. Currently, only the TencentDB for Redis Community Edition (excluding Redis 2.8) supports account settings.	Managing Account

September 2019

Update	Description	Documentation
Released Monitoring 2.0	Redis Monitoring 2.0 has been released, introducing over 16 new monitoring metrics, encompassing network latency, response errors, and other key performance indicators.	Monitoring features

August 2019

Update	Description	Documentation
Supported password-free access	You can submit a ticket to enable password-free access. We recommend that you restrict server access by using security groups after this feature is enabled.	Password-Free Access
Supported disabling high-risk commands online	The use of certain Redis commands may lead to service instability or inadvertent data deletion. Therefore, Tencent Cloud's Redis database provides the functionality to disable certain commands. The commands that can be disabled include flushall, flushdb, keys, hgetall, eval, evalsha, and script.	Disabling Commands
Supported Cluster Edition in DTS	DTS enables one-click migration of self-built Redis Cluster and Codis to the cloud, supporting 3.0, 3.2, and 4.0 Cluster version migration, as well as 2.8 and 3.2 Codis version migration.	Migration with DTS

©2013-2026 Tencent Cloud. All rights reserved.

July 2019

Update	Description	Documentation
Supported version upgrade with DTS	TencentDB for Redis instance version upgrades are facilitated through Data Transmission Service (DTS) using hot migration, ensuring uninterrupted service during the upgrade process and enabling real-time incremental data updates.	Version Upgrade with DTS
Redis 4.0 Standard Edition Released	Redis 4.0 Standard Edition is released and supports 1-master 5-replica mode, read/write separation, scaling without disconnection, and ultra-high availability.	Memory Edition (Standard Architecture)

October 2018

Update	Description	Documentation
Redis 4.0 Cluster Edition Released	Redis 4.0 Cluster Edition is released with 4 TB of storage capacity, can carry tens of millions of concurrent access requests, and supports scaling without storage capacity loss and automatic read/write separation.	Memory Edition (Cluster Architecture)

©2013-2026 Tencent Cloud. All rights reserved.

Announcements
Announcement of Full Decommissioning of TencentCloud API 2.0

Last updated：2024-02-26 14:36:06

Considering the higher latency and complexity of using the Cloud API 2.0 version, technical support for TencentDB for Redis API 2.0 will be discontinued, and the service will be fully decommissioned starting from **January 1, 2023** Beijing Time.
The current TencentDB for Redis API 2.0 has three interfaces that have not been decommissioned and are still being used by existing customers. We recommend that you upgrade these interfaces to the API 3.0 version as soon as possible to avoid any impact on your business.

List of APIs switched from v2.0 to v3.0

API 2.0	API 3.0	API Meaning
DescribeRedis	DescribeInstances	Querying instance list
GetBackupDownloadUrl	DescribeBackupUrl	The API is used to query the download URL of backup files.
GetRedisBackupList	DescribeInstanceBackups	Query the instance backup list

Note:
When using API 3.0, we recommend visiting Tencent Cloud SDK Center to obtain SDKs for various programming languages that support API 3.0. API Explorer can directly generate SDK call code, making it easier for you to implement code invocation.

©2013-2026 Tencent Cloud. All rights reserved.

Classic Network Will Be Disused

Last updated: 2024-02-26 14:37:17

Dear Tencent Cloud users:
To provide you with better services, Tencent Cloud has ceased the creation of all classic network products on **March 31, 2022**. The classic network products will be officially discontinued on **December 31, 2022**, and the corresponding services will be provided by Virtual Private Cloud (VPC) products.
As the mainstream cloud network environment, Virtual Private Cloud (VPC) covers all the functionalities of classic networks, offering a more secure, flexible, and stable network environment and user experience. If your business is still using classic networks, we recommend migrating to Virtual Private Cloud (VPC) to ensure a better service experience.

> **Note:**
> To better understand this network migration, you can refer to Migration Notes to address your concerns, and then follow the Migration Plan for implementation.

Contact Us

If you are unsure about the classic network instances under your current account or have any questions regarding service migration, please feel free to contact us. We are here to serve you wholeheartedly!

©2013-2026 Tencent Cloud. All rights reserved.

API for Querying Instance Big Key Will Be Disused

Last updated: 2024-02-26 14:32:27

Dear Tencent Cloud users:
TencentDB for Redis plans to discontinue the Big Key query interface DescribeInstanceMonitorBigKey on **October 31, 2022**. The corresponding DBbrain interface DescribeRedisTopBigKeys will provide the capability to query the list of big keys.
DBbrain is a product used for intelligent database diagnosis and optimization. It offers real-time performance diagnostics and security protection for users, assisting in efficiently identifying fault causes, providing optimization suggestions, and aiding in preventive measures from the source. With AI-powered parameter tuning capabilities, DBbrain enhances the overall performance of the database.
The DBbrain corresponding interface DescribeRedisTopBigKeys allows specifying the sorting field and Key type. This helps quickly obtain big Key information in instances, such as big Key expiration time, memory size, number of elements, and maximum element length.

List of deprecated APIs

`DescribeInstanceMonitorBigKey` API and its associated APIs. `escribeInstanceMonitorBigKeySizeDist` · and `DescribeInstanceMonitorBigKeyTypeDist` will be deprecated simultaneously.

API Name	Feature
DescribeInstanceMonitorBigKey	Queries big keys of an instance
DescribeInstanceMonitorBigKeySizeDist	Queries the big key size distribution of an instance
DescribeInstanceMonitorBigKeyTypeDist	Queries the big key type distribution of an instance

©2013-2026 Tencent Cloud. All rights reserved.

Announcement of TencentDB for Redis 1-Minute Granularity Discontinuation

Last updated: 2026-05-12 15:40:59

Cloud Monitor plans to **discontinue the alarm policies and QCE/REDIS namespace for TencentDB for Redis with a 1-minute granularity on September 30, 2022.** The following impacts will occur after the discontinuation:

- You can't receive the alarm notification of TencentDB for Redis Memory Edition (1-minute granularity).
- Data can't be displayed on the dashboard of TencentDB for Redis Memory Edition (1-minute granularity).
- You can't pull the monitoring data of TencentDB for Redis Memory Edition (1-minute granularity).

Migration Plan

- Alarm policy: You need to migrate the metrics with 1-minute monitoring granularity to those with 5-second monitoring granularity in TencentDB for Redis. Click to view the mappings of metrics.
- Dashboard: You need to migrate the dashboard of the existing TencentDB for Redis instances with 1-minute monitoring granularity to 5-second monitoring granularity. Click to view the mappings of metrics.
- API retrieval of monitoring metric data: You need to switch the QCE/REDIS namespace to QCE/REDIS_MEM and update the metric names accordingly. Click to view the mappings of metrics.

Because the names of metric with 1-minute monitoring granularity of TencentDB for Redis are not completely the same as the metric names in Mappings of Metrics. This document describes the mappings of monitoring metricswith 1-minute monitoring granularity and 5-second monitoring granularity of TencentDB for Redis.

Monitoring Metric Mappings Between 1-Minute and 5-Second Alarm Policy

1-Minute Metric	1-Minute Metric Parameter	5-Second Metric	5-Second Metric Parameter	Unit	MetricsDescription

©2013-2026 Tencent Cloud. All rights reserved.

CPU load	cpu_us_min	CPU utilization	cpu_util	%	Avg CPU usage
Maximum CPU load	CpuMaxUsMin	Max CPU utilization of a node	cpu_max_util	%	The maximum CPU utilization of a node (shard or replica) in an instance
MEM Usage	storage_min	MEM Usage	mem_used	MB	Memory capacity actually used, including data and cache
Capacity usage	storage_us_min	Memory utilization	mem_util	%	The ratio of the memory actually used to the total memory requested
Maximum capacity utilization	StorageMaxUsMin	Max memory utilization of a node	mem_max_util	%	The maximum memory utilization of a node (shard or replica) in an instance
Private Inbound Traffic	in_flow_min	Inbound Traffic	in_flow	Mb/s	Private Inbound Traffic
Inbound traffic utilization	in_flow_us_min	Inbound traffic utilization	in_bandwidth_util	%	The ratio of the actually used private inbound traffic to the maximum traffic
Private Outbound Traffic	out_flow_min	Traffic Out	out_flow	Mb/s	Private Outbound Traffic
Outbound traffic utilization	out_flow_us_min	Outbound traffic utilization	out_bandwidth_util	%	The ratio of the actually used private outbound traffic to the maximum traffic

©2013–2026 Tencent Cloud. All rights reserved.

Number of connections	connections_min	Connection Count	connections	Connections	Number of TCP connections to an instance
Connection utilization	connections_us_min	Connection utilization	connections_util	%	The ratio of the number of TCP connections to the maximum number of connections
Number of slow queries	slow_query_min	Slow query	cmd_slow	Threads	The number of command executions with a latency greater than the `slowlog-log-slower-than` configuration
Keys	keys_min	Keys	keys	Connections	Total number of keys stored in an instance (first-level keys)
Expired keys	expired_keys_min	Expired keys	expired	Connections	The number of keys expired in a time window, which is equal to the value of `expired_keys` output by the `info` command
Evicted keys	evicted_keys_min	Evicted keys	evicted	Conn	The number of keys evicted in a time window, which is equal to the value

©2013–2026 Tencent Cloud. All rights reserved.

				ections	of `evicted_keys` output by the `info` command
Average execution latency	latency_min	Average execution latency	latency_avg	ms	Average execution latency from Proxy to Redis server
Average read latency	latency_get_min	Average read latency	latency_read	ms	Average execution latency of read commands from proxy to Redis server
Average write latency	latency_set_min	Average write latency	latency_write	ms	Average execution latency of write command from proxy to Redis server
Average latency of other commands	latency_other_min	Average latency of other commands	latency_other	ms	The average execution latency of commands other than read and write commands from Proxy to Redis server
qps	qps_min	Total requests	commands	Queries/second	QPS, that is, the number of command executions per second
Read Requests	stat_get_min	Read Requests	cmd_read	Queries/sec	The number of read command executions per second

©2013–2026 Tencent Cloud. All rights reserved.

				o n d	
Write Request	stat_set_min	Write Request	cmd_write	Queries/second	The number of write command executions per second
Other Requests	stat_other_min	Other Requests	cmd_other	Queries/second	The number of command (excluding write and read commands) executions per second.
Big value requests	big_value_min	Big value requests	cmd_big_value	Queries/second	The number of executions of commands larger than 32 KB per second
Read request hits	stat_success_min	Read request hits	cmd_hits	Threads	The number of keys successfully requested by read commands, which is equal to the value of the `ke`

©2013–2026 Tencent Cloud. All rights reserved.

					`yspace_hits` metric output by the `info` command
Read request misses	stat_missed_min	Read request misses	cmd_miss	Threads	The number of keys unsuccessfully requested by read commands, which is equal to the value of the `keyspace_misses` metric output by the `info` command
Execution Errors	cmd_err_min	Execution Errors	cmd_err	Threads	Number of command execution errors, such as when a command does not exist or a parameter is incorrect
Read request hit rate	cache_hit_ratio_min	Read request hit rate	cmd_hits_ratio	%	Key hits/(key hits + key misses). This metric can reflect the situation of cache miss. When the access request quantity is 0, the value of this metric will be null

Mappings of Monitoring Metrics Between 1-Minute and 5-Second Dashboard

1-Minute Metric	1-Minute Metric Parameter	5-Second Metric	5-Second Metric Parameter	Unit	MetricsDescription
CPU utilization	cpu_us_min	Avg CPU usage	cpu_util	%	Avg CPU usage
Max CPU utilization of a shard	cpu_max_us_min	Max CPU utilization of a node	cpu_max_util	%	The maximum CPU utilization of a node (shard or replica) in an instance

©2013-2026 Tencent Cloud. All rights reserved.

MEM Usage	storage_min	MEM Usage	mem_used	MB	Memory capacity actually used, including data and cache
Memory utilization	storage_us_min	Memory utilization	mem_util	%	The ratio of the memory actually used to the total memory requested
Max Memory Utilization of Shard	storage_max_us_min	Max memory utilization of a node	mem_max_util	%	The maximum memory utilization of a node (shard or replica) in an instance
Inbound Traffic	in_flow_min	Inbound Traffic	in_flow	Mb/s	Private Inbound Traffic
Inbound traffic utilization	in_flow_us_min	Inbound traffic utilization	in_bandwidth_util	%	The ratio of the actually used private inbound traffic to the maximum traffic
Traffic Out	out_flow_min	Traffic Out	out_flow	Mb/s	Private Outbound Traffic
Outbound traffic utilization	out_flow_us_min	Outbound traffic utilization	out_bandwidth_util	%	The ratio of the actually used private outbound traffic to the maximum traffic
Connection Count	connections_min	Connection Count	connections	Connections	Number of TCP connections to an instance

©2013–2026 Tencent Cloud. All rights reserved.

Connection utilization	connections_us_min	Connection utilization	connections_util	%	The ratio of the number of TCP connections to the maximum number of connections
Slow query	slow_query_min	Slow query	cmd_slow	Threads	The number of command executions with a latency greater than the `slowlog-log-slower-than` configuration
Keys	keys_min	Keys	keys	Connections	Total number of keys stored in an instance (first-level keys)
Expired keys	expired_keys_min	Expired keys	expired	Connections	The number of keys expired in a time window, which is equal to the value of `expired_keys` output by the `info` command
Evicted keys	evicted_keys_min	Evicted keys	evicted	Connections	The number of keys evicted in a time window, which is equal to the value of `evicted_keys` output by the `info` command
Average executi	latency_min	Average execution	latency_avg	ms	Average execution latency from Proxy to Redis server

©2013–2026 Tencent Cloud. All rights reserved.

on latency		latency			
Average read latency	latency_get_min	Average read latency	latency_read	ms	Average execution latency of read commands from proxy to Redis server
Average write latency	latency_set_min	Average write latency	latency_write	ms	Average execution latency of write command from proxy to Redis server
Average latency of other commands	latency_other_min	Average latency of other commands	latency_other	ms	The average execution latency of commands other than read and write commands from Proxy to Redis server
qps	qps_min	Total requests	commands	Queries/second	QPS, that is, the number of command executions per second
Read Requests	stat_get_min	Read Requests	cmd_read	Queries/second	The number of read command executions per second

©2013–2026 Tencent Cloud. All rights reserved.

Number of write requests	stat_set_min	Write Request	cmd_write	Queries/second	The number of write command executions per second
Number of Get requests	cmdstat_get_min	Other Requests	cmd_other	Queries/second	The number of command (excluding write and read commands) executions per second.
Number of Getbit requests	cmdstat_getbit_min				
Number of Getrange requests	cmdstat_getrange_min				
Number of Mget requests	cmdstat_hget_min				
Number of Hgetall requests	cmdstat_hmget_min				
Number of Hmget	cmdstat_hmget_min				

©2013–2026 Tencent Cloud. All rights reserved.

requests	
Number of Hmset requests	cmdstat_hmset_min
Number of Hset requests	cmdstat_hset_min
Number of Hsetnx requests	cmdstat_hsetnx_min
Number of lset requests	cmdstat_lset_min
Number of Mget requests	cmdstat_mget_min
Number of Mset requests	cmdstat_mset_min
Number of Msetnx requests	cmdstat_msetnx_min
Number of Set	cmdstat_set_min

©2013–2026 Tencent Cloud. All rights reserved.

requests					
Number of Setbit requests	cmdstat_setbit_min				
Number of Setex requests	cmdstat_setex_min				
Number of Setnx requests	cmdstat_setnx_min				
Number of Setrange requests	cmdstat_setnx_min				
Big value requests	big_value_min	Big value requests	cmd_big_value	Queries/second	The number of executions of commands larger than 32 KB per second
Read request hits	stat_success_min	Read request hits	cmd_hits	Threads	The number of keys successfully requested by read commands, which is equal to the value of the `key space_hits` metric output by the `info` command

©2013–2026 Tencent Cloud. All rights reserved.

Read request misses	stat_missed_min	Read request misses	cmd_miss	Threads	The number of keys unsuccessfully requested by read commands, which is equal to the value of the `keyspace_misses` metric output by the `info` command
Execution Errors	cmd_err_min	Execution Errors	cmd_err	Threads	Number of command execution errors, such as when a command does not exist or a parameter is incorrect
Read request hit rate	cache_hit_ratio_min	Read request hit rate	cmd_hits_ratio	%	Key hits/(key hits + key misses). This metric can reflect the situation of cache miss. When the access request quantity is 0, the value of this metric will be null

Mappings between 1-minute and 5-second API monitoring metrics

Namespace

- 1-minute: Namespace=QCE/REDIS
- 5-second: Namespace=QCE/REDIS_MEM

Monitored metrics

Instance dimension (standard architecture)

1-Minute Metric	1-Minute Metric Parameter	5-Second Metric	5-Second Metric Parameter	Unit	MetricsDescription
CPU utilization	CpuUsMin	CPU utilization	CpuUtil	%	Avg CPU usage
MEM Usage	StorageMin	MEM Usage	MemUsed	MB	Memory capacity actually used, including data and cache

©2013-2026 Tencent Cloud. All rights reserved.

Memory utilization	StorageUsMin	Memory utilization	MemUtil	%	The ratio of the memory actually used to the total memory requested
Keys	KeysMin	Keys	Keys	Connections	Total number of keys stored in an instance (first-level keys)
Expired keys	ExpiredKeysMin	Expired keys	Expired	Connections	The number of keys expired in a time window, which is equal to the value of `expired_keys` output by the `info` command
Evicted keys	EvictedKeysMin	Evicted keys	Evicted	Connections	The number of keys evicted in a time window, which is equal to the value of `evicted_keys` output by the `info` command
Connection Count	ConnectionsMin	Connection Count	Connections	Connections	Number of TCP connections to an instance

©2013-2026 Tencent Cloud. All rights reserved.

Connection utilization	ConnectionsUsMin	Connection utilization	ConnectionsUtil	%	The ratio of the number of TCP connections to the maximum number of connections
Inbound Traffic	InFlowMin	Inbound Traffic	InFlow	Mb/s	Private Inbound Traffic
Inbound traffic utilization	InFlowUsMin	Inbound traffic utilization	InBandwidthUtil	%	The ratio of the actually used private inbound traffic to the maximum traffic
TrafficOut	OutFlowMin	Traffic Out	OutFlow	Mb/s	Private Outbound Traffic
Outbound traffic utilization	OutFlowUsMin	Outbound traffic utilization	OutBandwidthUtil	%	The ratio of the actually used private outbound traffic to the maximum traffic
Average execution latency	LatencyMin	Average execution latency	LatencyAvg	ms	Average execution latency from Proxy to Redis server
Average read latency	LatencyGetMin	Average read latency	LatencyRead	ms	Average execution latency of read commands from proxy to Redis server
Average write latency	LatencySetMin	Average write latency	LatencyWrite	ms	Average execution latency of write command from proxy to Redis server
Average latency of other commands	LatencyOtherMin	Average latency of other commands	LatencyOther	ms	The average execution latency of commands other than read and write commands from Proxy to Redis server

©2013–2026 Tencent Cloud. All rights reserved.

Total requests	QpsMin	Total requests	Commands	Queries/second	QPS, that is, the number of command executions per second
Read Requests	StatGetMin	Read Requests	CmdRead	Queries/second	The number of read command executions per second
Write Request	StatSetMin	Write Request	CmdWrite	Queries/second	The number of write command executions per second
Other Requests	StatOtherMin	Other Requests	CmdOther	Queries/seco	The number of command (excluding write and read commands) executions per second.

©2013–2026 Tencent Cloud. All rights reserved.

				nd	
Big value requests	BigValueMin	Big value requests	CmdBigValue	Queries/second	The number of executions of commands larger than 32 KB per second
Slow query	SlowQueryMin	Slow query	CmdSlow	Threads	The number of command executions with a latency greater than the `slowlog-log-slower-than` configuration
Read request hits	StatSuccessMin	Read request hits	CmdHits	Threads	The number of keys successfully requested by read commands, which is equal to the value of the `keyspace_hits` metric output by the `info` command
Read request misses	StatMissedMin	Read request misses	CmdMiss	Threads	The number of keys unsuccessfully requested by read commands, which is equal to the value of the `keyspace_misses` metric output by the `info` command
Execution Errors	CmdErrMin	Execution Errors	CmdErr	Threads	Number of command execution errors, such as when a command does not exist or a parameter is incorrect

©2013–2026 Tencent Cloud. All rights reserved.

Read request hit rate	CacheHitRatioMin	Read request hit rate	CmdHitsRatio	%	Key hits/(key hits + key misses). This metric can reflect the situation of cache miss. When the access request quantity is 0, the value of this metric will be null

Instance dimension (cluster architecture)

1-Minute Metric	1-Minute Metric Parameter	5-Second Metric	5-Second Metric Parameter	Unit	MetricsDescription
Avg CPU usage	CpuUsMin	CPU utilization	CpuUtil	%	Avg CPU usage
Max CPU utilization of a shard	CpuMaxUsMin	Max CPU utilization of a node	CpuMaxUtil	%	Highest CPU utilization value of all shards in a cluster
MEM Usage	StorageMin	MEM Usage	MemUsed	MB	Memory capacity actually used, including data and cache
Memory utilization	StorageUsMin	Memory utilization	MemUtil	%	The ratio of the memory actually used to the total memory requested
Max Memory Utilization of Shard	StorageMaxUsMin	Max memory utilization of a node	MemMaxUtil	%	The highest memory utilization value of all shards in a cluster
Keys	KeysMin	Keys	Keys	Connne	Total number of keys stored in an instance (first-level keys)

©2013-2026 Tencent Cloud. All rights reserved.

				ctions	
Expired keys	ExpiredKeysMin	Expired keys	Expired	Connections	The number of keys expired in a time window, which is equal to the value of `expired_keys` output by the `info` command
Evicted keys	EvictedKeysMin	Evicted keys	Evicted	Connections	The number of keys evicted in a time window, which is equal to the value of `evicted_keys` output by the `info` command
Connection Count	ConnectionsMin	Connection Count	Connections	Connections	Number of TCP connections to an instance
Connection utilization	ConnectionsUsMin	Connection utilization	ConnectionsUtil	%	The ratio of the number of TCP connections to the maximum number of connections
Inbound Traffic	InFlowMin	Inbound Traffic	InFlow	Mb/s	Private Inbound Traffic

©2013–2026 Tencent Cloud. All rights reserved.

Inbound traffic utilization	InFlowUsMin	Inbound traffic utilization	InBandwidthUtil	%	The ratio of the actually used private inbound traffic to the maximum traffic
TrafficOut	OutFlowMin	Traffic Out	OutFlow	Mb/s	Private Outbound Traffic
Outbound traffic utilization	OutFlowUsMin	Outbound traffic utilization	OutBandwidthUtil	%	The ratio of the actually used private outbound traffic to the maximum traffic
Average execution latency	LatencyMin	Average execution latency	LatencyAvg	ms	Average execution latency from Proxy to Redis server
Average read latency	LatencyGetMin	Average read latency	LatencyRead	ms	Average execution latency of read commands from proxy to Redis server
Average write latency	LatencySetMin	Average write latency	LatencyWrite	ms	Average execution latency of write command from proxy to Redis server
Average latency of other commands	LatencyOtherMin	Average latency of other commands	LatencyOther	ms	The average execution latency of commands other than read and write commands from Proxy to Redis server
Total requests	QpsMin	Total requests	Commands	Quer	QPS, that is, the number of command executions per second

©2013–2026 Tencent Cloud. All rights reserved.

				ies/second	
Read Requests	StatGetMin	Read Requests	CmdRead	Queries/second	The number of read command executions per second
Write Request	StatSetMin	Write Request	CmdWrite	Queries/second	The number of write command executions per second
Other Requests	StatOtherMin	Other Requests	CmdOther	Queries/second	The number of command (excluding write and read commands) executions per second.

©2013–2026 Tencent Cloud. All rights reserved.

Big value requests	BigValueMin	Big value requests	CmdBigValue	Queries/second	The number of executions of commands larger than 32 KB per second
Slow query	SlowQueryMin	Slow query	CmdSlow	Threads	The number of command executions with a latency greater than the `slowlog-log-slower-than` configuration
Read request hits	StatSuccessMin	Read request hits	CmdHits	Threads	The number of keys successfully requested by read commands, which is equal to the value of the `keyspace_hits` metric output by the `info` command
Read request misses	StatMissedMin	Read request misses	CmdMiss	Threads	The number of keys unsuccessfully requested by read commands, which is equal to the value of the `keyspace_misses` metric output by the `info` command
Execution Errors	CmdErrMin	Execution Errors	CmdErr	Threads	Number of command execution errors, such as when a command does not exist or a parameter is incorrect
Read request hit rate	CacheHitRatioMin	Read reque	CmdHitsRatio	%	Key hits/(key hits + key misses). This metric can reflect the situation of

©2013–2026 Tencent Cloud. All rights reserved.

		st hit rate			cache miss. When the access request quantity is 0, the value of this metric will be null

Cluster sharding

1-Minute Metric	1-Minute Metric Parameter	5-Second Metric	5-Second Metric Parameter	Unit	MetricsDescription
CPU utilization	CpuUsNodeMin	CPU utilization	CpuUtilNode	%	Avg CPU usage
MEM Usage	StorageNodeMin	MEM Usage	MemUsedNode	MB	Memory capacity actually used, including data and cache
Memory utilization	StorageUsNodeMin	Memory utilization	MemUtilNode	%	The ratio of the memory actually used to the total memory requested
Keys	KeysNodeMin	Keys	KeysNode	Connections	Total number of keys stored in an instance (first-level keys)
Expired keys	ExpiredKeysNodeMin	Expired keys	ExpiredNode	Connections	The number of keys expired in a time window, which is equal to the value of `expired_keys` output by the `info` command
Evicted keys	EvictedKeysNodeMin	Evicted keys	EvictedNode	Co	The number of keys evicted in a time window,

©2013-2026 Tencent Cloud. All rights reserved.

				nnections	which is equal to the value of `evicted_keys` output by the `info` command
Total requests	QpsNodeMin	Total requests	CommandsNode	Queries/second	QPS, that is, the number of command executions per second
Read Requests	StatGetNodeMin	Read Requests	CmdReadNode	Queries/second	The number of read command executions per second
Write Request	StatSetNodeMin	Write Request	CmdWriteNode	Queries/second	The number of write command executions per second
Other Request	StatOtherNodeMin	Other Reques	CmdOtherNode	Qu	The number of command (excluding write and read

©2013–2026 Tencent Cloud. All rights reserved.

s		ts		eries/second	commands) executions per second.
Slow query	SlowQueryNodeMin	Slow query	CmdSlowNode	Threads	The number of command executions with a latency greater than the `slowlog-log-slower-than` configuration
Read request hits	StatSuccessNodeMin	Read request hits	CmdHitsNode	Threads	The number of keys successfully requested by read commands, which is equal to the value of the `keyspace_hits` metric output by the `info` command
Read request misses	StatMissedNodeMin	Read request misses	CmdMissNode	Threads	The number of keys unsuccessfully requested by read commands, which is equal to the value of the `keyspace_misses` metric output by the `info` command
Execution Errors	CmdErrNodeMin	Execution Errors	CmdErr	Threads	Number of command execution errors, such as when a command does not exist or a parameter is incorrect
Read request hit rate	CacheHitRatioNodeMin	Read request hit rate	CmdHitsRatioNode	%	Key hits/(key hits + key misses). This metric can reflect the situation of cache miss. When the access request quantity is

©2013–2026 Tencent Cloud. All rights reserved.

					0, the value of this metric will be null

Description of parameters in each dimension

Parameter (5-Second Granularity)	Dimension Name	Dimension Explanation	Format	Mapping
Instances.N.Dimensions.0.Name	instanceid	Dimension name of the instance ID	Enter a string-type dimension name: instanceid.	Unchanged
Instances.N.Dimensions.0.Value	instanceid	Specific instance ID	Enter the specific Redis instance ID, for example: tdsql-123456, or the instance serial number, for example: crs-ifmymj41. You can query through the Redis Instance List Query Interface.	Unchanged
Instances.N.Dimensions.1.Name	rnodeid	Dimension name of the Redis node ID	Enter a string-type dimension name: rnodeid.	Corresponding 1-minute dimension name: clusterid.
Instances.N.Dimensions.1.Value	rnodeid	Specific Redis node ID	Enter a Redis node ID, which can be obtained through the DescribeInstanceNodeInfo API.	Corresponding 1-minute dimension name: clusterid.

©2013-2026 Tencent Cloud. All rights reserved.