

TencentDB for Redis Release Notes and Announcements 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

STencent Cloud

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

Release Notes and Announcements

Release Notes

Announcements

Classic Network Will Be Disused

API for Querying Instance Big Key Will Be Disused

1-Minute Granularity Will Be Disused

Monitoring Upgrade and Alarm Policy Changes

Release Notes and Announcements Release Notes

Last updated : 2024-04-22 15:53:44

March 2024

Update	Description	Document
Added global replication monitoring metrics.	Supports graphical monitoring of the command count deviation during the global replication process between the primary instance and disaster recovery instances, as well as the synchronization latency in remote synchronization.	Monitoring Metrics

February 2024

Update	Description	Document
One-click connectivity checker	If the connection to TencentDB for Redis instance fails, it is recommended to first use the one-click connectivity checker to identify the cause.	One-Click Connectivity Checker

January 2024

Update	Description	Document
Added monitoring metrics.	Adds a monitoring metric for the number of keys with expiration time set.	Monitoring feature (5-second granularity)

February 2023

Update	Description	Document
Supported the configuration of backup file download rules	To prevent database backup data from being dragged, TencentDB for Redis allows you to download backup files from specific servers on the allowed private network.	
Supported the selection of billing methods for clone instances	When cloning a complete new instance based on the current backup file, you can choose the pay-as-you-go billing mode.	

November 2022

Update	Description	Document
Released Redis 6.2	Redis 6.2 is fully released, which supports standard and cluster architectures, read/write separation, scaling without disconnection, and 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 easily view them locally.	Setting Instance Parameters

October 2022

Update	Description	Document
Optimized the instance list in the console	The Deployment Mode column is added in 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 .	Viewing 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
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	Document
Optimized statistical algorithm for traffic and throttling that triggers monitoring metrics	TencentDB for Redis now provides optimized statistical algorithm for inbound/outbound traffic and throttling.	_
Supported the 256 MB specification for the standard architecture	The minimum memory specification (256 MB) is available for TencentDB for Redis 4.0 and 5.0 Memory Edition (Standard Architecture) instances in more AZs.	Performance



June 2022

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

May 2022

Update	Description	Document
Supported minor version upgrade	TencentDB for Redis supports minor version upgrade, so that you can upgrade your instance to a newer version for more features.	Upgrading Instance Version
Supported proxy upgrade	Minor versions of the TencentDB for Redis proxy are released from time to time to add more database features or fix known bugs.	Upgrading Proxy

April 2022

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

deployed instances		
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	Document
Supported global replication	The global replication feature allows you to add and remove instances, switch instance roles, and delete replication groups.	Overview
Supported custom backup retention period	You can apply for customizing the retention period of backup files through an allowlist.	Backing up Data
Added the API for failure simulation	Instances in the same AZ support calling the KillMasterGroup API for failure simulation testing.	KillMasterGroup
Added the API for master- replica switch	Instances in the same AZ support calling the ChangeReplicaToMaster API to promote a replica node (group) to a master node for master-replica switch.	ChangeReplicaToMaster

December 2021

Update	Description	Document
Supported the 256 MB memory	The minimum memory specification of TencentDB for Redis 4.0 and 5.0 Memory Edition (Standard Architecture) instances can be 256 MB.	Performance



specification for	Currently, this is supported only in Shanghai Zone 5, Beijing Zone 6, and Guangzhou Zone 6.	
TencentDB		
for Redis 4.0		
and 5.0		
Memory		
Edition		
(Standard		
Architecture)		

November 2021

Update	Description	Document
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 allows you to enable public network access in the console, so that you can connect to your TencentDB for Redis instance from a CVM instance in different VPCs over the public network.	Configuring Public Network Address

July 2021

Update	Description	Document
Supported parameter template	Besides the system parameter templates provided by TencentDB for Redis, you can create custom parameter templates to configure parameters in batches as needed.	Applying Parameter Templates

June 2021

Update	Description	Document
Supported	TencentDB for Redis supports the automatic failover feature for proxy	Failover

automatic failover	nodes and Redis server nodes to ensure service availability.	
Auto-Failback	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.	Automatic Failback
Manually Promoting to Master Node (Group)	For multi-AZ deployed TencentDB for Redis instances in the standard/cluster architecture, you can manually promote a replica node/node group to master node/node group. You can deploy the master node in a specified AZ/node group.	Manually Promoting to Master Node (Group)
Reading Local Nodes Only	To reduce the access latency of a multi-AZ deployed instance, TencentDB for Redis allows you to read local nodes only.	Reading Local Nodes Only

March 2021

Update	Description	Document
Fully supported monitoring metrics at a 5-second granularity	TencentDB for Redis now supports the 5-second monitoring granularity. After the monitoring granularity of an instance is adjusted to five seconds, the monitoring metrics, proxy, and alarm policies of the instance, and the method of viewing monitoring data in the TCOP console will change.	Monitoring Upgrade and Alarm Policy Changes

December 2020

Update	Description	Document
Renamed TencentDB for Redis Hybrid Storage Edition TencentDB for Tendis	TencentDB for Redis Hybrid Storage Edition has been renamed TencentDB for Tendis. You can access TencentDB for Tendis in its own console.	TencentDB for Tendis

Supported	You can now deploy TencentDB for Redis master and replica nodes in	Multi-AZ
multi-AZ	different availability zones (AZs) of the same region. Multi-AZ deployed	Deployment
deployment	instances have higher availability and better disaster recovery capability	
	than single-AZ deployed instances.	

September 2020

Update	Description	Document
Supported the 5-second granularity for monitoring data	The monitoring feature of TencentDB for Redis has been updated, with the monitoring granularity being narrowed down from one minute to five seconds.	Monitoring at Five-Second Granularity

July 2020

Update	Description	Document
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
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	Document
Launched TencentDB for Redis Hybrid Storage Edition	TencentDB for Redis Hybrid Storage Edition is launched. With 100% compatibility with Redis protocols, it reduces memory costs by up to 80% and delivers a hot data performance comparable with that of Memory Edition, achieving a perfect balance among compatibility, performance, and costs.	Hybrid Storage Edition (Cluster Architecture)



April 2020

Update	Description	Document
Supported Redis 5.0 in DTS	You can now migrate data and upgrade instances to Redis 5.0 with DTS.	Migration with DTS Version Upgrade with DTS

March 2020

Update	Description	Document
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, read-only replica failover, and read-only replica unavailability.	Configuring Alarm
Supported adjusting bandwidth in the console	You can now adjust instance network bandwidth on the instance details page in the Redis console. The bandwidth specifications for Cluster Edition have been fully upgraded, with shard specifications upgraded to 384 Mbps.	-
Updated the monitoring view for Redis 2.8	To provide more accurate monitoring information, Tencent Cloud has updated the monitoring view of Redis 2.8 instance to cluster view. If you use APIs to get monitoring data from TCOP, you need to change the view parameter from redisuuid to cluster in the code.	-

February 2020

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



January 2020

Update	Description	Document
Released Redis 5.0	Redis 5.0 Standard and Cluster Editions have been released with all features of Redis 4.0 reserved, and support the latest STREAM data structure, new ZSET commands ZPOPMIN and ZPOPMAX, scaling without disconnection, 4 TB of storage capacity, and tens of millions of concurrent QPS.	Memory Edition (Standard Architecture) Memory Edition (Cluster Architecture)

October 2019

Update	Description	Document
Supported access management	You can now create policies through access management which grant sub-accounts permissions to use the resources they need.	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	Document
Released monitoring 2.0	Redis monitoring 2.0 has been released, adding more than 16 monitoring metrics, including network delay, response error and other monitoring metrics.	Monitoring

August 2019

Update	Description	Document
Supported	To enable password-free access, submit a ticket. We recommend that	Password-Free



password-free access	you limit server access by using security groups after this feature is enabled.	Access
Supported disabling high- risk commands online	TencentDB for Redis is now able to disable some Redis commands that may cause service instability or accidentally delete data, including flushall, flushdb, keys, hgetall, eval, evalsha, and script.	Disabling Commands
Supported Cluster Edition in DTS	You can now migrate self-created Redis Cluster (3.0, 3.2, and 4.0) or Codis (2.8 and 3.2) databases to TencentDB for Redis with DTS.	Migration with DTS

July 2019

Update	Description	Document
Supported version upgrade with DTS	TencentDB for Redis instance version can now be upgraded through DTS with hot migration, which guarantees instance service continuity during the upgrade process and can update incremental data in real time.	Version Upgrade with DTS
Released Redis 4.0 Standard Edition	Redis 4.0 Standard Edition supports 1-master 5-replica mode, read/write separation, scaling without disconnection, and ultra high availability.	Memory Edition (Standard Architecture)

October 2018

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

Announcements Classic Network Will Be Disused

Last updated : 2022-11-08 17:18:57

Dear user,

Tencent Cloud stopped the support of the creation of any new resources in the classic network on **March 31, 2022** and will officially discontinue the classic network product from **December 31, 2022**. After that, all classic network-based resources will be migrated to VPC.

As the current mainstream cloud network environment, VPC has all the features of the classic network while providing a more secure, flexible, and stable network environment and user experience. If your business is still using the classic network, we recommend you migrate it to VPC for a better service experience.

Note:

For more information, see Migration Solutions

Contact Us

For any questions, contact us.

API for Querying Instance Big Key Will Be Disused

Last updated : 2022-11-17 11:00:02

Dear user,

TencentDB for Redis will discontinue API DescribeInstanceMonitorBigKey for querying the big key of an instance on October 31, 2022, but this feature will continue to serve through the API DescribeRedisTopBigKeys of DBbrain.

TencentDB for DBbrain (DBbrain) is an intelligent database diagnosis and optimization product that provides real-time performance diagnosis and security protection services. It troubleshoots efficiently, offers solutions to database exceptions, and helps you prevent exceptions at the source. It can also help improve the overall database performance with its AI-powered parameter tuning capabilities.

DescribeRedisTopBigKeys of DBbrain can specify the sorting field and type of the Key to quickly query of big key information of an instance, including expiration time, memory size, element quantity, and the max element length.

List of deprecated APIs

DescribeInstanceMonitorBigKey API and its associated APIs.

```
escribeInstanceMonitorBigKeySizeDist •and DescribeInstanceMonitorBigKeyTypeDist will be deprecated simultaneously.
```

APIName	Description
DescribeInstanceMonitorBigKey	Queries the big key of an instance
DescribeInstanceMonitorBigKeySizeDist	Queries the big key size distribution of an instance
DescribeInstanceMonitorBigKeyTypeDist	Queries the big key type distribution of an instance

1-Minute Granularity Will Be Disused

Last updated : 2023-03-29 15:36:39

CM will discontinue the 1-minute granularity alarm policy for TencentDB for Redis and QCE/REDIS namespace, which will have the following impacts:

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 Schemes

Alarm policy: You need to migrate the alarm policy of the existing TencentDB for Redis instances with 1-minute monitoring granularity to those with 5-second monitoring granularity. 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.

Pulling monitoring metric data through APIs: You need to switch QCE/REDIS to QCE/REDIS_MEM, and modify the names of pulled metrics simultaneously.

Because 1-minute monitoring granularity of TencentDB for Redis is not completely the same as the metric names in Mappings of Metrics. This document describes the mappings of monitoring metrics of TencentDB for Redis 1-minute monitoring granularity and 5-second monitoring granularity for you to compare and migrate.

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	Description
CPU Load	cpu_us_min	CPU Utilization	cpu_util	%	Average CPU utilization
Max CPU Load	CpuMaxUsMin	Max CPU Utilization of Node	cpu_max_util	%	The maximum Cl utilization of a no (shard or replica) instance
Used Memory	storage_min	Used Memory	mem_used	MB	Actually used me capacity, includir capacity for data cache
Capacity	storage_us_min	Memory	mem_util	%	The ratio of the a



Utilization		Utilization			used memory to requested total m
Max Capacity Utilization	StorageMaxUsMin	Max Memory Utilization of Node	mem_max_util	%	The maximum ar all node (shard o replica) memory utilizations in an instance
Private Network Inbound Traffic	in_flow_min	Inbound Traffic	in_flow	MB/s	Private network inbound traffic
Inbound Traffic Utilization	in_flow_us_min	Inbound Traffic Utilization	in_bandwidth_util	%	The ratio of the a used private inbc traffic to the maxi traffic
Private Network Outbound Traffic	out_flow_min	Outbound Traffic	out_flow	MB/s	Private network outbound traffic
Outbound Traffic Utilization	out_flow_us_min	Outbound Traffic Utilization	out_bandwidth_util	%	The ratio of the a used private outk traffic to the maxi traffic
Connections	connections_min	Connections	connections	-	The number of T connections to ar instance
Connection Utilization	connections_us_min	Connection Utilization	connections_util	%	The ratio of the n of TCP connectic the maximum nu of connections
Slow Logs	slow_query_min	Slow Logs	cmd_slow	-	The number of command execut with a latency gre than the slowl log-slower-t configuration
Total Keys	keys_min	Total Keys	keys	-	The total number keys (level-1 key



					instance storage
Expired Keys	expired_keys_min	Expired Keys	expired	-	The number of keepired in a time window, which is to the value of expired_key output by the in command.
Evicted Keys	evicted_keys_min	Evicted Keys	evicted	-	The number of ker evicted in a time window, which is to the value of evicted_key output by the in command.
Average Execution Latency	latency_min	Average Execution Latency	latency_avg	ms	The average exe latency between proxy and the Re server
Average Read Latency	latency_get_min	Average Read Latency	latency_read	ms	The average exe latency of read commands betwo the proxy and the server
Average Write Latency	latency_set_min	Average Write Latency	latency_write	ms	The average exe latency of write command betwe proxy and the Re server
Average Latency of Other Commands	latency_other_min	Average Latency of Other Commands	latency_other	ms	The average exe latency of comma (excluding write a read commands) between the pro- the Redis server
QPS	qps_min	Total Requests	commands	Counts/sec	Queries per secc
Read Requests	stat_get_min	Read Requests	cmd_read	Counts/sec	The number of re command execut



					per second
Write Requests	stat_set_min	Write Requests	cmd_write	Counts/sec	The number of w command execut per second
Other Requests	stat_other_min	Other Requests	cmd_other	Counts/sec	The number of command (exclu- write and read commands) exec per second
Big Value Requests	big_value_min	Big Value Requests	cmd_big_value	Counts/sec	The number of executions of req larger than 32 KE second
Read Request Hits	stat_success_min	Read Request Hits	cmd_hits	-	The number of kersuccessfully required comman which is equal to value of the keyspace_hi metric output by info comman
Read Request Misses	stat_missed_min	Read Request Misses	cmd_miss	-	The number of ke unsuccessfully requested by rea commands, whic equal to the value keyspace_mi metric output by info comman
Execution Errors	cmd_err_min	Execution Errors	cmd_err	-	The number of command execut errors. For example the command do exist, parameters incorrect, etc.
Read Request Hit Rate	cache_hit_ratio_min	Read Request Hit Rate	cmd_hits_ratio	%	Key hits/(key hits misses). This me can reflect the sit of cache miss. W

		the access reque quantity is 0, the of this metric will
		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	Description
CPU Utilization	cpu_us_min	Average CPU utilization	cpu_util	%	Average CPU utilization
Max CPU Utilization of Shard	cpu_max_us_min	Max CPU Utilization of Node	cpu_max_util	%	The maximum utilization of a (shard or replic instance
Used Memory	storage_min	Used Memory	mem_used	MB	Actually used capacity, inclu capacity for da cache
Memory Utilization	storage_us_min	Memory Utilization	mem_util	%	The ratio of the used memory requested tota
Max Memory Utilization of Shard	storage_max_us_min	Max Memory Utilization of Node	mem_max_util	%	The maximum all node (sharc replica) memo utilizations in a instance
Inbound Traffic	in_flow_min	Inbound Traffic	in_flow	MB/s	Private networ inbound traffic
Inbound Traffic Utilization	in_flow_us_min	Inbound Traffic Utilization	in_bandwidth_util	%	The ratio of the used private ir traffic to the ma traffic
Outbound	out_flow_min	Outbound	out_flow	MB/s	Private networ

Traffic		Traffic			outbound traffi
Outbound Traffic Utilization	out_flow_us_min	Outbound Traffic Utilization	out_bandwidth_util	%	The ratio of the used private o traffic to the ma traffic
Connections	connections_min	Connections	connections	-	The number of connections to instance
Connection Utilization	connections_us_min	Connection Utilization	connections_util	%	The ratio of the of TCP connect the maximum i of connections
Slow Logs	slow_query_min	Slow Logs	cmd_slow	-	The number of command exe with a latency than the slo log-slower- configuration
Total Keys	keys_min	Total Keys	keys	-	The total numł keys (level-1 k instance stora
Expired Keys	expired_keys_min	Expired Keys	expired	-	The number of expired in a tin window, which to the value of expired_k output by the command.
Evicted Keys	evicted_keys_min	Evicted Keys	evicted	-	The number of evicted in a tin window, which to the value of evicted_ke output by the command.
Average Execution Latency	latency_min	Average Execution Latency	latency_avg	ms	The average e latency betwee proxy and the server



Average Read Latency	latency_get_min	Average Read Latency	latency_read	ms	The average e latency of reac commands be the proxy and server
Average Write Latency	latency_set_min	Average Write Latency	latency_write	ms	The average e latency of write command betv proxy and the server
Average Latency of Other Commands	latency_other_min	Average Latency of Other Commands	latency_other	ms	The average e latency of com (excluding writ read command between the p the Redis serv
QPS	qps_min	Total Requests	commands	Counts/sec	Queries per se
Read Requests	stat_get_min	Read Requests	cmd_read	Counts/sec	The number of command exempts a command exempts
Write Requests	stat_set_min	Write Requests	cmd_write	Counts/sec	The number of command exemption per second
get Requests	cmdstat_get_min	Other Requests	cmd_other	Counts/sec	The number of command (exc
getbit Requests	cmdstat_getbit_min				write and read commands) e> per second
getrange Requests	cmdstat_getrange_min				
hget Requests	cmdstat_hget_min				
hgetall Requests	cmdstat_hmget_min				
hmget Requests	cmdstat_hmget_min				



hmset Requests	cmdstat_hmset_min				
hset Requests	cmdstat_hset_min				
hsetnx Requests	cmdstat_hsetnx_min				
lset Requests	cmdstat_lset_min				
mget Requests	cmdstat_mget_min				
mset Requests	cmdstat_mset_min				
msetnx Requests	cmdstat_msetnx_min				
set Requests	cmdstat_set_min				
setbit Requests	cmdstat_setbit_min				
setex Requests	cmdstat_setex_min				
setnx Requests	cmdstat_setnx_min				
setrange Requests	cmdstat_setnx_min				
Big Value Requests	big_value_min	Big Value Requests	cmd_big_value	Counts/sec	The number of executions of r larger than 32 second
Read Request Hits	stat_success_min	Read Request Hits	cmd_hits	-	The number of successfully re by read comm which is equal value of the keyspace_1



					metric output k
Read Request Misses	stat_missed_min	Read Request Misses	cmd_miss	-	The number of unsuccessfully requested by r commands, wl equal to the va keyspace_r metric output k info comn
Execution Errors	cmd_err_min	Execution Errors	cmd_err	-	The number of command exe- errors. For exa the command exist, paramet incorrect, etc.
Read Request Hit Rate	cache_hit_ratio_min	Read Request Hit Rate	cmd_hits_ratio	%	Key hits/(key h misses). This r can reflect the of cache miss. the access rec quantity is 0, th of this metric v null.

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

Namespace

1-minute: Namespace=QCE/REDIS

5-second: Namespace=QCE/REDIS_MEM

Monitoring metrics

Instance dimension (standard architecture)

1-Minute Metric	1-Minute Metric Parameter	5-Second Metric	5-Second Metric Parameter	Unit	Description
CPU Utilization	CpuUsMin	CPU Utilization	CpuUtil	%	Average CPU utilization
Used	StorageMin	Used	MemUsed	MB	Actually used memo



Memory		Memory			capacity, including t capacity for data an cache
Memory Utilization	StorageUsMin	Memory Utilization	MemUtil	%	The ratio of the actu used memory to the requested total men
Total Keys	KeysMin	Total Keys	Keys	-	The total number of keys (level-1 keys) i instance storage
Expired Keys	ExpiredKeysMin	Expired Keys	Expired	-	The number of keys expired in a time window, which is eq to the value of expired_keys output by the info command.
Evicted Keys	EvictedKeysMin	Evicted Keys	Evicted	-	The number of keys evicted in a time window, which is eq to the value of evicted_keys output by the info command.
Connections	ConnectionsMin	Connections	Connections	-	The number of TCP connections to an instance
Connection Utilization	ConnectionsUsMin	Connection Utilization	ConnectionsUtil	%	The ratio of the num of TCP connections the maximum numb of connections
Inbound Traffic	InFlowMin	Inbound Traffic	InFlow	MB/s	Private network inbound traffic
Inbound Traffic Utilization	InFlowUsMin	Inbound Traffic Utilization	InBandwidthUtil	%	The ratio of the actu used private inboun traffic to the maximu traffic
Outbound Traffic	OutFlowMin	Outbound Traffic	OutFlow	MB/s	Private network outbound traffic



Outbound Traffic Utilization	OutFlowUsMin	Outbound Traffic Utilization	OutBandwidthUtil	%	The ratio of the actu used private outbou traffic to the maximu traffic
Average Execution Latency	LatencyMin	Average Execution Latency	LatencyAvg	LatencyAvg ms	
Average Read Latency	LatencyGetMin	Average Read Latency	LatencyRead	ms	The average execut latency of read commands betweer the proxy and the Re server
Average Write Latency	LatencySetMin	Average Write Latency	LatencyWrite	ms	The average execut latency of write command between proxy and the Redis server
Average Latency of Other Commands	LatencyOtherMin	Average Latency of Other Commands	LatencyOther	ms	The average execut latency of commanc (excluding write and read commands) between the proxy a the Redis server
Total Requests	QpsMin	Total Requests	Commands	Counts/sec	Queries per second
Read Requests	StatGetMin	Read Requests	CmdRead	Counts/sec	The number of read command executior per second
Write Requests	StatSetMin	Write Requests	CmdWrite	Counts/sec	The number of write command executior per second
Other Requests	StatOtherMin	Other Requests	CmdOther	Counts/sec	The number of command (excluding write and read commands) execution per second
Big Value	BigValueMin	Big Value	CmdBigValue	Counts/sec	The number of



Requests		Requests			executions of reque larger than 32 KB pr second
Slow Logs	SlowQueryMin	Slow Logs	CmdSlow	-	The number of command execution with a latency greate than the slowlog log-slower-than configuration
Read Request Hits	StatSuccessMin	Read Request Hits	CmdHits	-	The number of keys successfully reques 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	-	The number of keys unsuccessfully requested by read commands, which is equal to the value of keyspace_miss metric output by the info command.
Execution Errors	CmdErrMin	Execution Errors	CmdErr	-	The number of command executior errors. For example the command does exist, parameters ar incorrect, etc.
Read Request Hit Rate	CacheHitRatioMin	Read Request Hit Rate	CmdHitsRatio	%	Key hits/(key hits + misses). This metric can reflect the situat of cache miss. Whe the access request quantity is 0, the val 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	Description
Average CPU utilization	CpuUsMin	CPU Utilization	CpuUtil	%	Average CPU utilization
Max CPU Utilization of Shard	CpuMaxUsMin	Max CPU Utilization of Node	CpuMaxUtil	%	The highest CPU utilization value of a shards in a cluster
Used Memory	StorageMin	Used Memory	MemUsed	MB	Actually used memo capacity, including t capacity for data an cache
Memory Utilization	StorageUsMin	Memory Utilization	MemUtil	%	The ratio of the actu used memory to the requested total mem
Max Memory Utilization of Shard	StorageMaxUsMin	Max Memory Utilization of Node	MemMaxUtil	%	The highest memory utilization value of a shards in a cluster
Total Keys	KeysMin	Total Keys	Keys	-	The total number of keys (level-1 keys) i instance storage
Expired Keys	ExpiredKeysMin	Expired Keys	Expired	-	The number of keys expired in a time window, which is eq to the value of expired_keys output by the info command.
Evicted Keys	EvictedKeysMin	Evicted Keys	Evicted	-	The number of keys evicted in a time window, which is eq to the value of evicted_keys output by the info command.



Connections	ConnectionsMin	Connections	Connections	-	The number of TCP connections to an instance
Connection Utilization	ConnectionsUsMin	Connection Utilization	ConnectionsUtil	%	The ratio of the num of TCP connections the maximum numb of connections
Inbound Traffic	InFlowMin	Inbound Traffic	InFlow	MB/s	Private network inbound traffic
Inbound Traffic Utilization	InFlowUsMin	Inbound Traffic Utilization	InBandwidthUtil	%	The ratio of the actu used private inboun traffic to the maximu traffic
Outbound Traffic	OutFlowMin	Outbound Traffic	OutFlow	MB/s	Private network outbound traffic
Outbound Traffic Utilization	OutFlowUsMin	Outbound Traffic Utilization	OutBandwidthUtil	%	The ratio of the actu used private outbou traffic to the maximu traffic
Average Execution Latency	LatencyMin	Average Execution Latency	LatencyAvg	ms	The average execut latency between the proxy and the Redis server
Average Read Latency	LatencyGetMin	Average Read Latency	LatencyRead	ms	The average execut latency of read commands betweer the proxy and the Re server
Average Write Latency	LatencySetMin	Average Write Latency	LatencyWrite	ms	The average execut latency of write command between proxy and the Redis server
Average Latency of Other Commands	LatencyOtherMin	Average Latency of Other Commands	LatencyOther	ms	The average execut latency of commanc (excluding write anc read commands)



					between the proxy a the Redis server
Total Requests	QpsMin	Total Requests	Commands	Counts/sec	Queries per second
Read Requests	StatGetMin	Read Requests	CmdRead	Counts/sec	The number of read command executior per second
Write Requests	StatSetMin	Write Requests	CmdWrite	Counts/sec	The number of write command executior per second
Other Requests	StatOtherMin	Other Requests	CmdOther	Counts/sec	The number of command (excluding write and read commands) execution per second
Big Value Requests	BigValueMin	Big Value Requests	CmdBigValue	Counts/sec	The number of executions of requea larger than 32 KB pe second
Slow Logs	SlowQueryMin	Slow Logs	CmdSlow	-	The number of command execution with a latency greate than the slowlog log-slower-than configuration
Read Request Hits	StatSuccessMin	Read Request Hits	CmdHits	-	The number of keys successfully reques 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	-	The number of keys unsuccessfully requested by read commands, which is equal to the value of keyspace_misse



					metric output by the info command.
Execution Errors	CmdErrMin	Execution Errors	CmdErr	-	The number of command executior errors. For example the command does exist, parameters ar incorrect, etc.
Read Request Hit Rate	CacheHitRatioMin	Read Request Hit Rate	CmdHitsRatio	%	Key hits/(key hits + misses). This metric can reflect the situal of cache miss. Whe the access request quantity is 0, the val of this metric will be null.

Cluster sharding

1-Minute Metric	1-Minute Metric Parameter	5-Second Metric	5-Second Metric Parameter	Unit	Description
CPU Utilization	CpuUsNodeMin	CPU Utilization	CpuUtilNode	%	Average CPU utilization
Used Memory	StorageNodeMin	Used Memory	MemUsedNode	MB	Actually used merr capacity, including capacity for data a cache
Memory Utilization	StorageUsNodeMin	Memory Utilization	MemUtilNode	%	The ratio of the act used memory to th requested total me
Total Keys	KeysNodeMin	Total Keys	KeysNode	-	The total number c keys (level-1 keys) instance storage
Expired Keys	ExpiredKeysNodeMin	Expired Keys	ExpiredNode	-	The number of key expired in a time window, which is e to the value of expired_keys



					output by the int command.
Evicted Keys	EvictedKeysNodeMin	Evicted Keys	EvictedNode	-	The number of key evicted in a time window, which is e to the value of evicted_keys output by the int command.
Total Requests	QpsNodeMin	Total Requests	CommandsNode	Counts/sec	Queries per secon
Read Requests	StatGetNodeMin	Read Requests	CmdReadNode	Counts/sec	The number of rea command executic per second
Write Requests	StatSetNodeMin	Write Requests	CmdWriteNode	Counts/sec	The number of writ command executic per second
Other Requests	StatOtherNodeMin	Other Requests	CmdOtherNode	Counts/sec	The number of command (excludi write and read commands) execu per second
Slow Logs	SlowQueryNodeMin	Slow Logs	CmdSlowNode	-	The number of command executic with a latency grea than the slowlo log-slower-tha configuration
Read Request Hits	StatSuccessNodeMin	Read Request Hits	CmdHitsNode	-	The number of key successfully reque by read commands which is equal to the value of the keyspace_hits metric output by th info command
Read Request Misses	StatMissedNodeMin	Read Request Misses	CmdMissNode	-	The number of key unsuccessfully requested by read



					commands, which equal to the value keyspace_miss metric output by th info command
Execution Errors	CmdErrNodeMin	Execution Errors	CmdErr	-	The number of command executic errors. For exampl the command does exist, parameters a incorrect, etc.
Read Request Hit Rate	CacheHitRatioNodeMin	Read Request Hit Rate	CmdHitsRatioNode	%	Key hits/(key hits 4 misses). This metr can reflect the situa of cache miss. Wh the access reques quantity is 0, the va of this metric will b null.

Description of parameters in each dimension

Parameter (5-Second Granularity)	Dimension Name	Description	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 a specific Redis instance ID, such as tdsql-123456, which can be queried through the DescribeInstances API and can also be an instance string such as crs-ifmymj41.	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 specific Redis node ID, which can be queried through the DescribeInstanceNodeInfo API.	Corresponding 1-minute dimension name: clusterid.
--------------------------------	---------	------------------------------	--	---

Monitoring Upgrade and Alarm Policy Changes

Last updated : 2023-05-23 11:00:17

Monitoring Granularity

TencentDB for Redis now supports one-minute and five-second monitoring granularities. Since October 2020, monitoring at the five-second granularity has been supported, providing more monitoring metrics and proxy monitoring data. For more information, see Update Notes of Monitoring at Five-Second Granularity.

Changes of monitoring granularity

One-minute monitoring granularity

Instances created before October 20, 2020 only support the one-minute monitoring granularity, but they are gradually upgraded to support the five-second granularity.

View the monitoring data in the TCOP console: TCOP console > TencentDB > **Redis (1-minute granularity)**.

Five-second monitoring granularity

Instances created after October 20, 2020 support both one-minute and five-second monitoring granularities. View the monitoring data in the TCOP console: TCOP console > **TencentDB** > **Redis (5-second granularity)**.

Notes of monitoring granularity upgrade

To support five-second monitoring granularity, the proxy of your TencentDB for Redis instances needs to be upgraded to the latest version.

Note that the proxy upgrade will cause a short disconnection. The business needs to reconnect to the proxy after the upgrade completes.

1. Upgrade by Tencent Cloud backend: Tencent Cloud is upgrading all instances to support five-second monitoring granularity. You will be notified via SMS, email, or Message Center before the upgrade starts.

2. Upgrade by yourself in the TencentDB console: you can soon manually upgrade instances in the console.

3. After all instances are upgraded, the one-minute monitoring granularity will be no longer supported.

Changes of monitoring metrics

After the monitoring granularity is narrowed from one minute to five seconds, monitoring metric names are changed and some new metrics are supported, as shown below:

1-minute metric	5-second metric	Description



CpuUsMin	CpuUtil	Average CPU utilization	
CpuMaxUs	CpuMaxUtil	The maximum CPU utilization of a node (shard or replica) in an instance	
StorageMin	MemUsed	Actually used memory capacity, including the capacity for data and cache	
StorageUsMin	MemUtil	The ratio of the actually used memory to the requested total memory	
StorageMaxUs	MemMaxUtil	The maximum among all node (shard or replica) memory utilizations in an instance	
KeysMin	Keys	The total number of keys (level-1 keys) in instance storage	
ExpiredKeysMin	Expired	The number of keys expired in a time window, which is equal to the value of expired_keys output by the info command.	
EvictedKeysMin	Evicted	The number of keys evicted in a time window, which is equal to the value of evicted_keys output by the info command.	
ConnectionsMin	Connections	The number of TCP connections to an instance	
ConnectionsUsMin	ConnectionsUtil	The ratio of the number of TCP connections to the maximum number of connections	
InFlowMin	InFlow	Private network inbound traffic	
InFlowUs	InBandwidthUtil	The ratio of the actually used private inbound traffic to the maximum traffic	
-	InFlowLimit	The number of times inbound traffic triggers a traffic limit	
OutFlowMin	OutFlow	Private network outbound traffic	
OutFlowUs	OutBandwidthUtil	The ratio of the actually used private outbound traffic to the maximum traffic	
-	OutFlowLimit	The number of times outbound traffic triggers a traffic limit	
LatencyMin	LatencyAvg	The average execution latency between the proxy and the Redis server	
-	LatencyMax	The maximum execution latency between the proxy and the Redis server	



-	LatencyP99	The P99 latency between the proxy and the Redis server	
LatencyGetMin	LatencyRead	The average execution latency of read commands between the proxy and the Redis server	
LatencySetMin	LatencyWrite	The average execution latency of write command between the proxy and the Redis server	
LatencyOtherMin	LatencyOther	The average execution latency of commands (excluding write and read commands) between the proxy and the Redis server	
QpsMin	Commands	Queries per second	
StatGetMin	CmdRead	The number of read command executions per second	
StatSetMin	CmdWrite	The number of write command executions per second	
StatOtherMin	CmdOther	The number of command (excluding write and read commands) executions per second	
BigValueMin	CmdBigValue	The number of executions of requests larger than 32 KB per second	
-	CmdKeyCount	The number of keys accessed by a command per second	
-	CmdMget	The number of MGET commands executed per second	
SlowQueryMin	CmdSlow	The number of command executions with a latency greaterthan the slowlog-log-slower-than configuration	
StatSuccessMin	CmdHits	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.	
StatMissedMin	CmdMiss	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.	
CmdErrMin	CmdErr	The number of command execution errors. For example, the command does not exist, parameters are incorrect, etc.	
CacheHitRatioMin	CmdHitsRatio	Key hits/(key hits + key misses). This metric can reflect the situation of cache miss.	

Viewing the monitoring granularity of an instance

Check the value of the InstanceSet.MonitorVersion field returned by the DescribeInstances API. If the value is 5s, this instance supports the monitoring granularity of five seconds; if the value is 1m, it supports only the monitoring granularity of one minute.

Log in to the TencentDB for Redis console, click an instance name to enter the instance management page, select System Monitoring > Monitoring Metrics, and click the Period drop-down list at the top. If you can select 5 seconds from the drop-down list, this instance supports the monitoring granularity of five seconds, or else it supports only the monitoring granularity of one minute.

Alarm Changes

Changes of alarm policy configurations

After the monitoring metrics are upgraded, you need to configure one-minute-granularity and five-second-granularity alarm policies in different windows in the TCOP console.

Basic Information	1	
Policy Name	Up to 20 characters	
Notes	Up to 100 characters. Only Chinese and English characters, numbers, underscores, and hyphens are allowed.	
Monitor Type	Cloud Product Monitoring	
Policy Type	redis 💌	
Project 🕢	Redis / CKV	
	Cloud Database / Redis / Memory Edition / Instance	s. The current account has dynamic threshold
Configure Alarm	Cloud Database / Redis / Memory Edition / Proxy	
ct of monitoring	Cloud Database / Redis / Memory Edition / Redis Node	

Impact of monitoring upgrade

After the monitoring granularity is narrowed from one minute to five seconds, you need to migrate the one-minutegranularity alarm policies to the five-second-granularity alarm policies. The monitoring metrics applicable to the fivesecond granularity alarm policies are different from those applicable to one-minute granularity alarm policies. For more information, see Changes of monitoring metrics.

After the monitoring granularity is narrowed down to five seconds:

Monitoring data at both one-minute and five-second granularities are reported temporarily, that is, Cloud Monitor will stop reporting the one-minute-granularity data in the future.

One-minute-granularity alarm policies are valid temporarily.

The default five-second-granularity alarm policy is associated. Please specify alarm recipients for the default policy.

Migrating alarm policies

Manual migration: copy the existing one-minute-granularity alarm policies as the five-second-granularity alarm policies, but you need to configure alarm recipients for the five-second-granularity alarm policies. Automatic migration: after the monitoring granularity upgrade completes, the existing one-minute-granularity alarm policies will be automatically migrated to the five-second-granularity alarm policies, and you will be notified via SMS, email, or Message Center.