


TencentDB for MongoDB

Product Introduction

Product Documentation



Copyright Notice

©2013-2019 Tencent Cloud. All rights reserved.

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

Trademark Notice



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

Service Statement

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

Contents

Product Introduction

- Product Overview

- Benefits

- Use Cases

- Region and Availability Zone

- Storage Engine

- Disaster Recovery Architecture

- Service Description

- Features

- Use Limit

- Related Protocols

 - Protocol Difference

Product Introduction

Product Overview

Last updated : 2019-09-29 14:46:50

TencentDB for MongoDB

TencentDB for MongoDB is high-performance distributed data storage service powered by Tencent Cloud that enables MongoDB, an open source NoSQL database. It is fully compatible with MongoDB protocols and is well suited for non-relational database use cases.

Note:

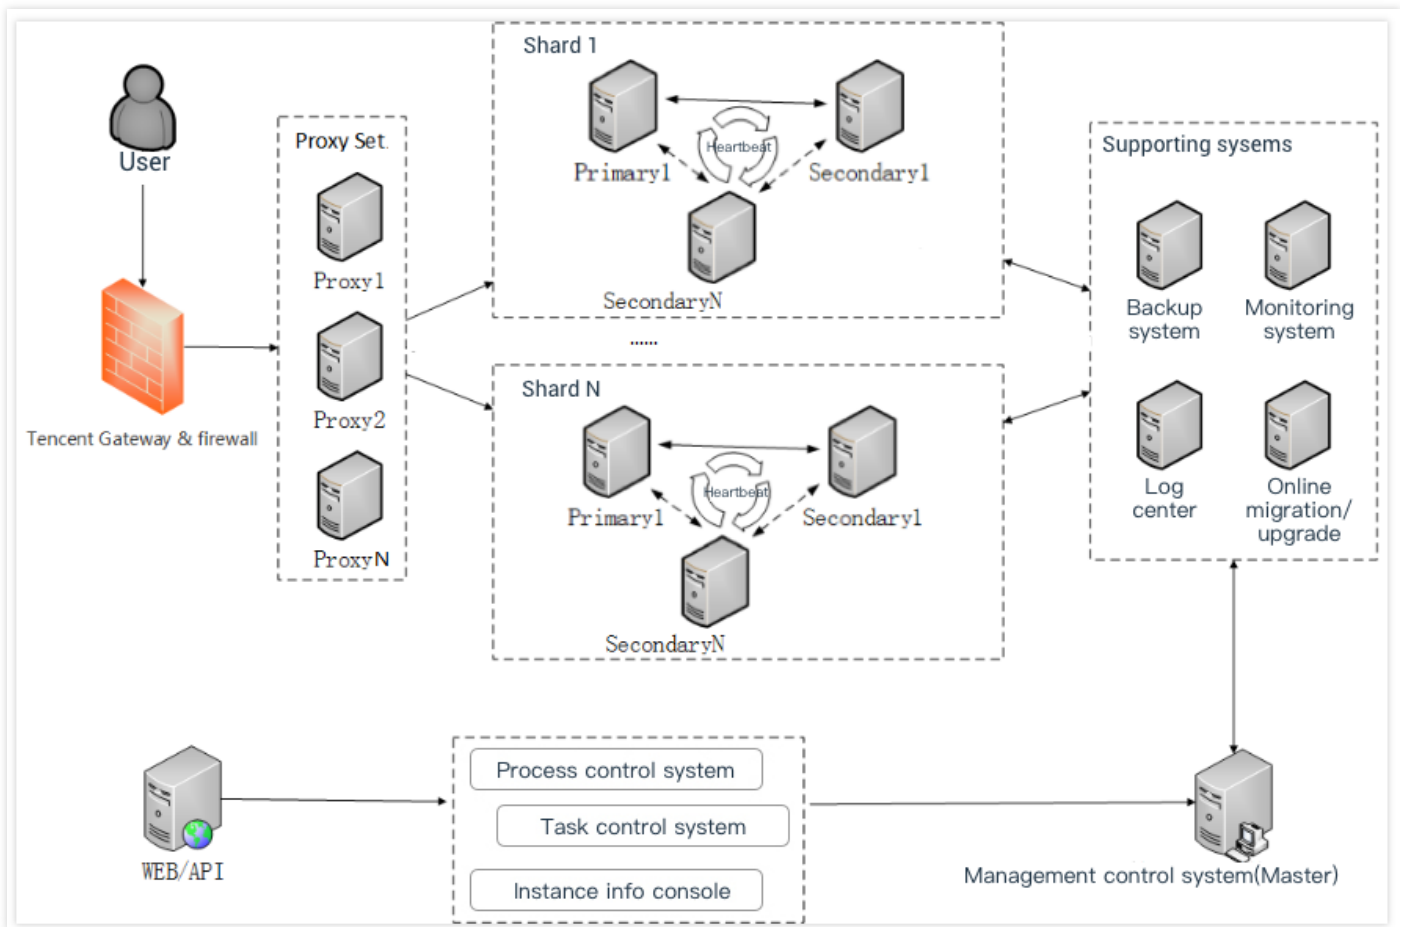
MongoDB® is the registered trademark of MongoDB, Inc. The product TencentDB for MongoDB is not authorized by MongoDB, Inc. and has no relation with MongoDB, Inc. Tencent's use of MongoDB software is based on the open source license.

Product Features

- TencentDB for MongoDB provides data storage services for Internet applications via Tencent Cloud platform.
- TencentDB for MongoDB is fully compatible with MongoDB protocols, and well suited for managing traditional relational data as well as cache and non-relational data. It supports MapReduce which allows parallel computation running on large-scale data sets.
- MongoDB cluster service provided by TencentDB for MongoDB is a reliable, easy-to-use, and high-performance service. Each instance has at least one master-slave replica set or a sharded cluster containing multiple replica sets.
- TencentDB for MongoDB has various functions such as backup and capacity expansion, ensuring data security and scalability.

Product Architecture

The following diagram illustrates the architecture of TencentDB for MongoDB system.



Benefits

Last updated : 2019-08-12 11:08:40

TencentDB for MongoDB Strengths

1. **Convenient:** It enables users to quickly set up MongoDB instances in Tencent Cloud environment, where you can directly access MongoDB instances through URI with no installments needed.
2. **Easy to use:** It is fully compatible with MongoDB protocol, so you can access instances via MongoDB protocol-based client and seamlessly migrate your existing MongoDB applications to the cloud platform.
3. **Secure:** It provides online backup (at least three copies) of your data to ensure data security. The data backup system also allows you to restore your data to several days ago for disaster recovery.
4. **High performance:** It deploys high-performance, high density, high capacity SSD storage servers for handling high-frequency access requests.
5. **Peace of mind:** It provides 24/7 professional, downtime-free, transparent service of server capacity expansion and data migration; enables comprehensive monitoring features for real-time service quality assessment.

TencentDB for MongoDB vs. Self-hosted MongoDB

TencentDB for MongoDB provides NoSQL service and is much easier to deploy, manage, and expand compared to on-premises MongoDB database. It is on demand and pay-as-you-go just like other cloud products, making it very cost effective. See below table for details.

Item	TencentDB for MongoDB	Self-hosted MongoDB
Cost	Zero hardware and software expenditure, multiple server options (high IO/high IO (10 GB)), and pay-as-you-use.	Hardware: high hardware cost - a single storage server is expensive (you need to buy 3 machines to build a highly available master-slave replication set, causing resource redundancy). Software: high labor cost - Professional database administrators are needed.
Service availability	99.95% service availability, industry-leading standards, 24/7 professional service, one-on-one instruction, and QQ remote assistance.	Need to solve failures and build master/slave or RAID by yourself.

Item	TencentDB for MongoDB	Self-hosted MongoDB
Reliability	99.9996% reliability, comprehensive auto data backup and lossless recovery mechanism, and real-time hot backup which allows you to restore data anytime within the last 5 days. Note: If the size of the data operated between two backups exceeds the oplog size, you cannot roll it back to a time between the two backups.	The reliability depends on the hardware failure rate and on the database management skills of technical personnel.
System security	Anti-DDoS attacks; repair various database and host security vulnerabilities in time.	Expensive to deploy yourself and you have to fix database security vulnerabilities by yourself.
Real-time monitoring	Multi-dimensional monitoring, and failure warning for ease of use.	You need to develop your own monitoring system, and OPS personnel are required to repair failures overnight.
Business scale-out	You can scale on-demand capacity through just one click with fast deployment.	You need to spend extra time on purchasing hardware, hosting data centers, and redeploying applications.
Resource utilization	Request on demand and 100% resource utilization to reduce your cost.	You have to spend on extra resources during peak periods, which lowers the average load and thus resource utilization rate.

TencentDB for MongoDB can resolve common self-hosted MongoDB operation issues, such as performance bottlenecks, OPS difficulties, data reliability, and availability problems:

- 1. Break performance bottleneck:** It adopts brand-new PCI-E SSD storage media and new generation storage engine; it provides customizable performance enhancing features to help improve your database performance.
- 2. Reduce OPS difficulty:** It features automatic monitor alarm for over 20 monitoring metrics and supports Big Data import/export and templated parameter modification to make it easy for business deployment.
- 3. High service availability:** It supports MongoDB master-slave with added hot backups for automatic disaster recovery and failover and the process is transparent to users. Like MongoDB, it allows you to read data from the slave first native MongoDB, which ensures high concurrent-read capability.

-
4. **High reliability:** It supports free data backup within 7 days, and enables private network firewall and public network DDoS protection.

Use Cases

Last updated : 2019-08-12 11:08:56

TencentDB is a cross-platform database. It is reliable, capable, and scalable, and can work with most of schemaless scenarios.

Below are some typical application scenarios.

Games

MongoDB is suitable for games managing their back-end databases, as it is more flexible in handling dynamic players' personal context, such as information about equipment and credit, through the use of embedded documents. It allows you to easily view, query and update the data, and the "schemaless" style saves you time and efforts adjusting database structures, significantly shortening the version iteration cycle.

MongoDB can also be used as a cache server to logically allocate hot data. It is comparable to other common cache servers in terms of performance while being more flexible in the way you query data.

Mobile

TencentDB for MongoDB supports two-dimensional spatial indexing, which makes it easy to query geo-location relationships and retrieve user geographic location data, thus achieving the implementation of geo-based map applications and such features as People/Places Nearby. MongoDB can also be used to store user information and user-generated content such as WeChat Moments.

Internet of Things (IoT)

MongoDB makes it easy for managing Tbs IoT data generated from user endpoint devices, such as medical sensors and GPS. MongoDB lets you store, process and multi-dimensionally analyze data (e.g., device log info) created by connected devices. The distributed TencentDB for MongoDB sharded clusters provide unlimited storage capacity and online storage expansion that allows you to easily process IoT Big Data.

Logistics

TencentDB for MongoDB can store the ever-changing order information in nested JSON so that you can view all the changes by running just one query.

Live Video Broadcasting (LVB)

TencentDB for MongoDB is built to storage volumes of data from virtual gifts, user chats, logs, etc, and to analyze the data through rich aggregate queries.

Region and Availability Zone

Last updated : 2019-08-14 20:42:24

Region

Tencent Cloud managed data centers are located around the world, including 4 regions in China (South, East, North, and Southwest), and also in Southeast Asia, Asia Pacific, West United States, East United States, North America, Europe, etc. We are deploying more nodes globally and enhancing network coverage.

You can create instances in the following regions.

- Guangzhou, Shanghai, Beijing, Chengdu, Chongqing, Hong Kong (China), Singapore, Seoul, Mumbai, Bangkok, Silicon Valley, Toronto, Virginia, Frankfurt, and Moscow.
- Available finance zones: Shenzhen Finance Zone and Shanghai Finance Zone.

While cloud products deployed within the same Region can communicate with each other over the private network, the private networks for the resources in different accounts are completely isolated from each other.

- Cloud products in different Regions cannot communicate with each other over the private network.
- Because each Region is independent, cross-region access to Tencent Cloud resources is not supported through private networks.
- When you purchase Tencent Cloud resources, we recommend you select a Region closest to your customers to minimize connection latency.

Availability Zone

In a Region, Availability Zones refer to Tencent Cloud physical IDCs with independent power facilities and networks. Availability Zones are designed to prevent single point failures (except for large-scale natural disasters or major power failures) from affecting other Availability Zones in the same Region to ensure your business availability. Availability Zones in the same Region are connected via low-latency private networks.

Storage Engine

Last updated : 2019-08-29 10:55:45

TencentDB for MongoDB supports both MongoRocks and WiredTiger engines.

MongoRocks Engine

MongoRocks is based on RocksDB, a KV store that is ideal for MongoDB document structures. Currently, it supports MongoDB 3.2 (stable version), MongoDB 3.4 (stable version), and MongoDB 3.6 (unstable version which does not support some features in v3.6 such as Change Stream).

Pros and Cons of MongoRocks

As MongoRocks is based on RocksDB, it has the common advantages and disadvantages of Log Structured Merge (LSM) storage systems, including:

- **Read amplification**

In RocksDB, reads are first made to MemTable. If no MemTable is found, RocksDB will scan through SST files layer by layer from the newest to the oldest (from top to bottom) until the desired data is located. This process may require more than one I/O. Although the read amplification can be reduced with Bloom filter and Cache, it is still quite obvious in scenarios such as range query.

- **Write amplification**

Each write in RocksDB needs to be persisted as level-0 SST files by Immutable MemTable, and the Level-0 files need to be compacted into level-1 files. Thus one write request will generate multiple disk I/Os. The lifespan of an SSD is subject to the number of writes to it, so special attention should be paid to write amplification for SSDs.

- **Disk amplification**

All writes are appends, so expired data will not be immediately cleaned up.

- **Write limitation**

Write Stall may cause lags in RocksDB. If the writing speed is higher than the speed of MemTable persistence and SST files compaction, the engine will limit the write speed; otherwise, SST files will increase dramatically in number, which will further degrade read performance. Therefore, lags may occur if the writing speed is too high.

Optimization of MongoRocks by Tencent Cloud

Background

In the native MongoRocks, when the size of the oplog reaches the upper limit, the delete oplog action will be triggered, which will generate many disk I/Os after write amplification. The oplog has the following features:

- Oplogs are created chronologically, so they are also sequenced by key when stored in RocksDB.
- There is an insert option but no update option, so SST files do not need compaction.
- Oplogs are deleted from the oldest to the newest, which means that the deletion starts with the earliest SST file.

Optimizations

Given the features of the oplog, we put oplogs and oplog metadata separately into two column families and clean up oplogs through compaction. In this way, only very few I/Os will be generated. Please find the specific process below:

- Determine whether extra oplogs need to be deleted by checking whether the space occupied exceeds the upper limit by the size of 2 SST files.
- Obtain the SST file metadata of the oplogs and scan through the SST files starting from the earliest one until the total size of the files scanned exceeds the space needed.
- Call the maximum deletion time of the compaction filter and set it as the maximum key for the SST file cluster mentioned above.
- Actively call CompactRange to delete the oplogs. With such optimization, only very few I/Os are generated for the oplog clean-up.

Applicable Scenarios of MongoRocks

Currently, MongoRocks is widely used in the industry. For example, cached cold data and WeChat billing cold data within Tencent are stored with the MongoRocks engine.

- **Scenarios that require minimal latency and glitches**

Practice shows that the performance of the WiredTiger engine is excellent when its cache size is larger than the index size; however, the database will trigger eviction if the index is not completely in the memory. In this case, the request delay will increase significantly. In contrast, MongoRocks has a much more stable performance.

- **Multi-table scenarios**

In the WiredTiger engine, each table and index is stored as a single file. When there are too many tables, a large number of small files will be generated on the disk, which will seriously compromise performance. Practical experience shows that the performance of the database will deteriorate when the number of collections reaches several thousand. In contrast, the MongoRocks engine does not create a file for each table, which avoids this problem.

- **Cold data storage scenario**

For cost considerations, cold data is usually stored on inexpensive storage media such as SATA. The performance of the WiredTiger engine on HDDs is not satisfactory either, especially in the case of eviction. By contrast, the MongoRocks engine converts random writes of I/O to sequential writes, which is very friendly to HDDs.

WiredTiger Engine

WiredTiger has a typical B-tree structure.

Engine Comparison

Test environment:

CPU: Intel Xeon 2.3 GHz, 24 cores

Memory: 50 GB Cache

Disk: PCIE-SSD

Version: MongoDB v3.2, WiredTiger engine, and MongoRocks engine

Test Performance	Test Result
Disk capacity consumption	<p>Test model: compare the size of disk data files after writing 500 GB of documents (each 500 B in size).</p> <p>Conclusion: for the same data, MongoRocks consumes slightly more disk capacity than WiredTiger.</p>
100% read performance comparison	<p>Test model: read data randomly in a single collection of 500 GB data, and after the data is warmed up for 10 minutes, record the QPS value within the next 10 minutes.</p> <p>Conclusion: MongoRocks is slightly better than WiredTiger in terms of read performance and has a more stable QPS in read-only scenarios.</p>
100% write performance comparison	<p>Test model: compare the writing speed within the 10 minutes after 500 GB of documents (each 500 B in size) are written.</p> <p>Conclusion: MongoRocks is significantly better than WiredTiger in terms of write performance, and MongoRocks does not show the performance deterioration WiredTiger shows during the writing process.</p>

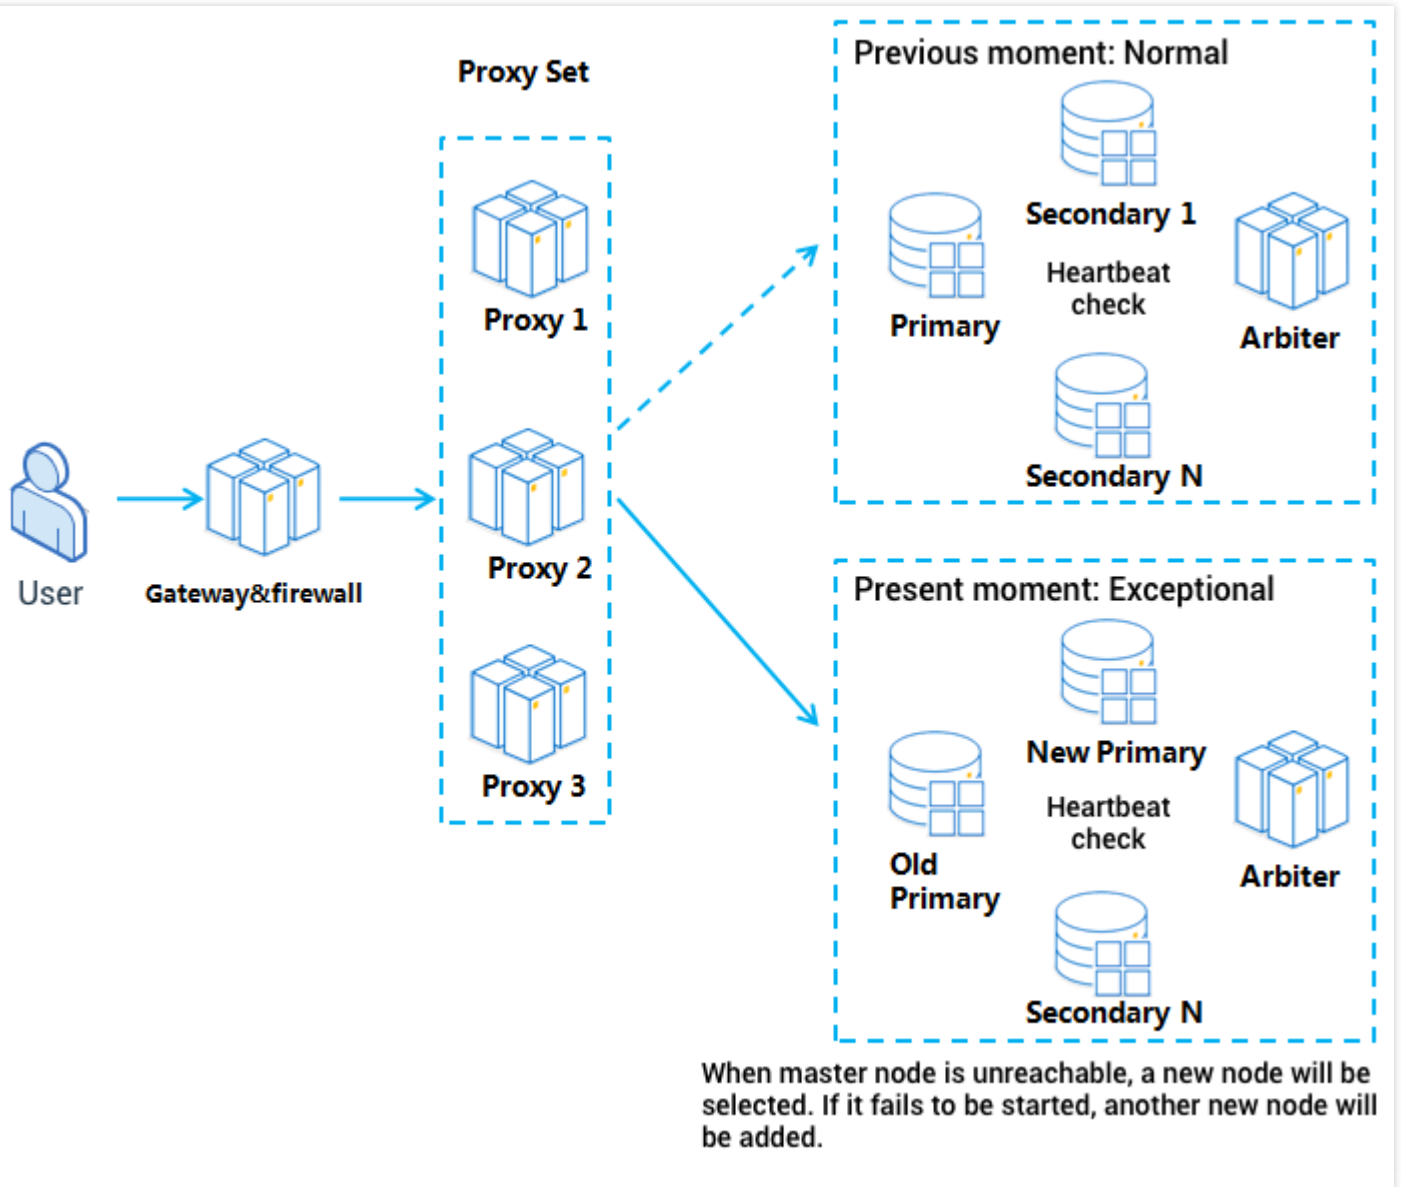
Test Performance	Test Result
1:1 read/write performance comparison	<p>Test model: read and write with a 1:1 ratio after 500 GB of documents (each 500 B in size) are written, after the data is warmed up for 10 minutes, record the QPS value within the next 10 minutes.</p> <p>Conclusion: MongoRocks and WiredTiger tie in performance in the 1:1 read/write scenario, but WiredTiger has some performance glitches.</p>

Disaster Recovery Architecture

Last updated : 2019-08-12 11:09:24

TencentDB for MongoDB adopts master/slave hot backup architecture. When the master node fails, the service will automatically switch to a slave node. Since a 10-second disconnection may occur during the master/slave switch, you need to set up the auto-restart. The following illustrates how auto disaster recovery works.

1. When the master node is unavailable due to an accident, the system will automatically select an alternative master node in the cluster.
2. If the master node fails, it will become a slave node after the restart. However, if the restart fails, a new node will be added to the cluster to maintain the specified cluster size.
3. Similarly, when a slave node is unavailable, the system will attempt to restart the node or add a new one.



Service Description

Last updated : 2018-09-14 15:27:10

1 Engine

Support WiredTiger engine and Rocks engine.

2 Cluster Types

There are two options when creating a MongoDB instance:

Node Configuration	Description
Replica set cluster	1 master, 2 slaves. Includes 1 Primary node and 2 Secondary nodes
Sharding cluster	A single instance contains at least 5 shards, and each shard consists of at least 3 nodes. It is expandable.

For replica set instances, you can set up a priority to read from Secondary first in your driver after connecting to the MongoDB service. For replica set instances, nodes can be expanded vertically. For sharding instances, nodes can be expanded vertically and shards can be expanded horizontally to support business scenarios of massive data.

3. Limit on Number of Connections

MEM specification	Max connections
2G	1500
4G	1500
6G	1500
8G	1500
12G	1500

MEM specification	Max connections
16G	2500
24G	3500
32G	4500
48G	6000
64G	9000
128G	15000
240G	15000
512G	15000

Note: The upper limit on number of connections is applicable to instances rather than nodes.

4 Connection User Name

There are two built-in default users: "rwuser" and "mongouser". The role for the built-in users is [readWriteAnyDatabase+dbAdmin](#), that is, the users can read and write any database, but are not permitted to perform critical operations.

Depending on the version of TencentDB for MongoDB, some instances only have rwuser (we will upgrade such instances and will contact you before doing so).

You can also use the Tencent Cloud MongoDB console for account authorization and permission management to meet requirements in various business scenarios.

5 Avoid Filling Up Disk

Write operation will be prohibited when the instance disk is 100% occupied, thus please expand capacity in time according to your business development. Contact customer service if you encounter such situation.

Features

Last updated : 2018-09-14 15:24:54

Performance Test Instructions

Tool

[Yahoo! Cloud Serving Benchmark](#)

Test Method

1. A single data entry is 1 KB
2. Occupy 80% of the instance capacity. For example: if the instance capacity is 100 GB, put 80 GB of data in it
3. Perform 50% read and 50% update operations to get QPS (Queries Per Second) data
4. Since there can be difference between the performances under different scenarios, please select instances with the appropriate specifications according to your business demand and test data

Performance Data of High IO Instances

CPU	memory	QPS
1 cores	2 G	3000
2 cores	4 G	5000
2 cores	6 G	6000
4 cores	8 G	9000
4 cores	12 G	14000
6 cores	16 G	20000
10 cores	24 G	25000
12 cores	32 G	27000
18 cores	48 G	30000
24 cores	64 G	33000

Performance Data of High IO Instance (10 Gigabyte)

CPU	Memory	Number of sets	Number of documents in a single set	Test data set (GB)	Runtime (S)	Data sampling interval (S)	Average QPS (rounded)
2 cores	4 G	2	100 million	79	1800	10	5000
4 cores	8 G	3	100 million	118	1800	10	9000
6 cores	16 G	6	100 million	237	1800	10	20000
12 cores	32 G	6	200 million	426	1800	10	27000
24 cores	64 G	15	200 million	1161	1800	10	33000
24 cores	128 G	30	200 million	2317	1800	10	36000
32 cores	240 G	40	200 million	3031	1800	10	39000

Use Limit

Last updated : 2019-09-26 20:51:15

1 Version & Engine

1) The current domestic regions including Beijing, Shanghai, and Guangzhou support MMAPv1 engine and WiredTiger engine

2 Replica Set & Sharding Cluster

2.1 Replica Set

There are two options when creating a MongoDB instance:

Node Configuration	Description
1 master, 2 slaves	Includes 1 Primary node and 2 Secondary nodes

There is a group of proxies (mongos in the MongoDB service component) before each replica set, and the URI for service connection appears like this:

```
mongodb://rwuser:password@10.66.77.88:27017/admin?authMechanism=MONGODB-CR
```

You can set up a priority to read from Secondary first in your driver after connecting to the MongoDB service, if required.

2.2 Sharding Cluster

[Sharding Introduction](#)

3 Connection User Name

There are two built-in default users: "rwuser" and "mongouser". The role for the built-in users is [readWriteAnyDatabase+dbAdmin](#), that is, the users can read and write any database, but are not permitted to perform critical operations.

Depending on the version of Tencent Cloud MongoDB, some instances only have rwuser (we will upgrade such instances and will contact you before doing so).

You can also use the Tencent Cloud MongoDB console for account and permission management to meet your business needs.

4 Avoid Filling Up Disk

Write operation will be prohibited when the instance disk is 100% occupied, therefore you need to expand capacity in time according to your business development. If you have such problem, you can contact customer service for help.

Related Protocols

Protocol Difference

Last updated : 2019-08-12 11:09:38

MongoDB and DynamoDB

MongoDB is a high-performance, non-relational, schema-free, document-oriented database that supports dynamic queries and full indexing. Performance excellence and the ease of use make MongoDB the most widely used NoSQL database for data management.

Core advantages of MongoDB:

1. Document Oriented Storage. Suitable for storing and managing collections of documents, where one collection can hold the unlimited size of documents.
2. Schema-free. Data is stored in the form of JSON style documents, which can have varying sets of fields, with different types for each field.
3. Index on any of attributes, including the internal objects. Geospatial-based indexes are also supported.
4. Rich queries. MongoDB supports a rich query language.
5. Aggregation capabilities, such as count and group.
6. Replication and data recovery. MongoDB's master-slave replication system provides a variety of features such as data backup, disaster recovery and read extension.

AWS DynamoDB is a fully managed proprietary NoSQL service that supports document and key-value data structures and provides fast, predictable performance at any scale.

Core advantages of DynamoDB:

1. Scalability. DynamoDB is built on a fully distributed and non-shared architecture that can scale the table assets to a large number of servers.
2. Predictable performance. The average latency of AWS DynamoDB server can be as low as a few milliseconds.
3. Easy to manage. As a fully managed service, AWS DynamoDB saves your time and money from the complexities of setting up and running a distributed database cluster with high reliability
4. Built-in fault tolerance. AWS DynamoDB's built-in fault tolerance supports automatic and synchronous replication of your data to multiple availability zones in a Region.
5. Flexibility. AWS DynamoDB does not have a fixed schema.
6. Strong consistency and atomic counter. AWS DynamoDB offers consistent read operations to ensure your data up-to-date. The atomic counter is supported.

Based on the existing NoSQL module framework, Tencent Cloud adds protocol-compatibility to AWS DynamoDB database service. Now DynamoDB is compatible with most basic protocols to meet database access requirements. Also, it supports instance-level backup and rollback, and automatic disaster recovery mechanism to ensure high reliability and availability of the service.

1. If you are a DynamoDB development enthusiast, Tencent Cloud database service is good for addressing your needs as a developer.
2. If you are from outside of China and need to deploy DynamoDB services in China for your projects that have been developed using the DynamoDB protocol, you can access Tencent Cloud Database via the DynamoDB protocol with minimum modification needed.