

TencentDB for MongoDB

Product Intro

Product Introduction



Copyright Notice

©2013-2018 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 Intro

Intro

Advantages

Scenarios

Region Selection

Service Description

Features

Use Limit

Product Intro

Intro

Last updated : 2018-09-14 15:24:18

1. About TencentDB for MongoDB

TencentDB for MongoDB is a high-performance distributed data storage service created by Tencent Cloud based on MongoDB, the world's most promising open source NoSQL database.

It is 100% compatible with MongoDB protocol, and is highly compatible with DynamoDB protocol. It is suitable for non-relational database-oriented scenarios such as games, logistics, social networking, Internet of Things, and videos.

2. Key Features

TencentDB for MongoDB has the following key features:

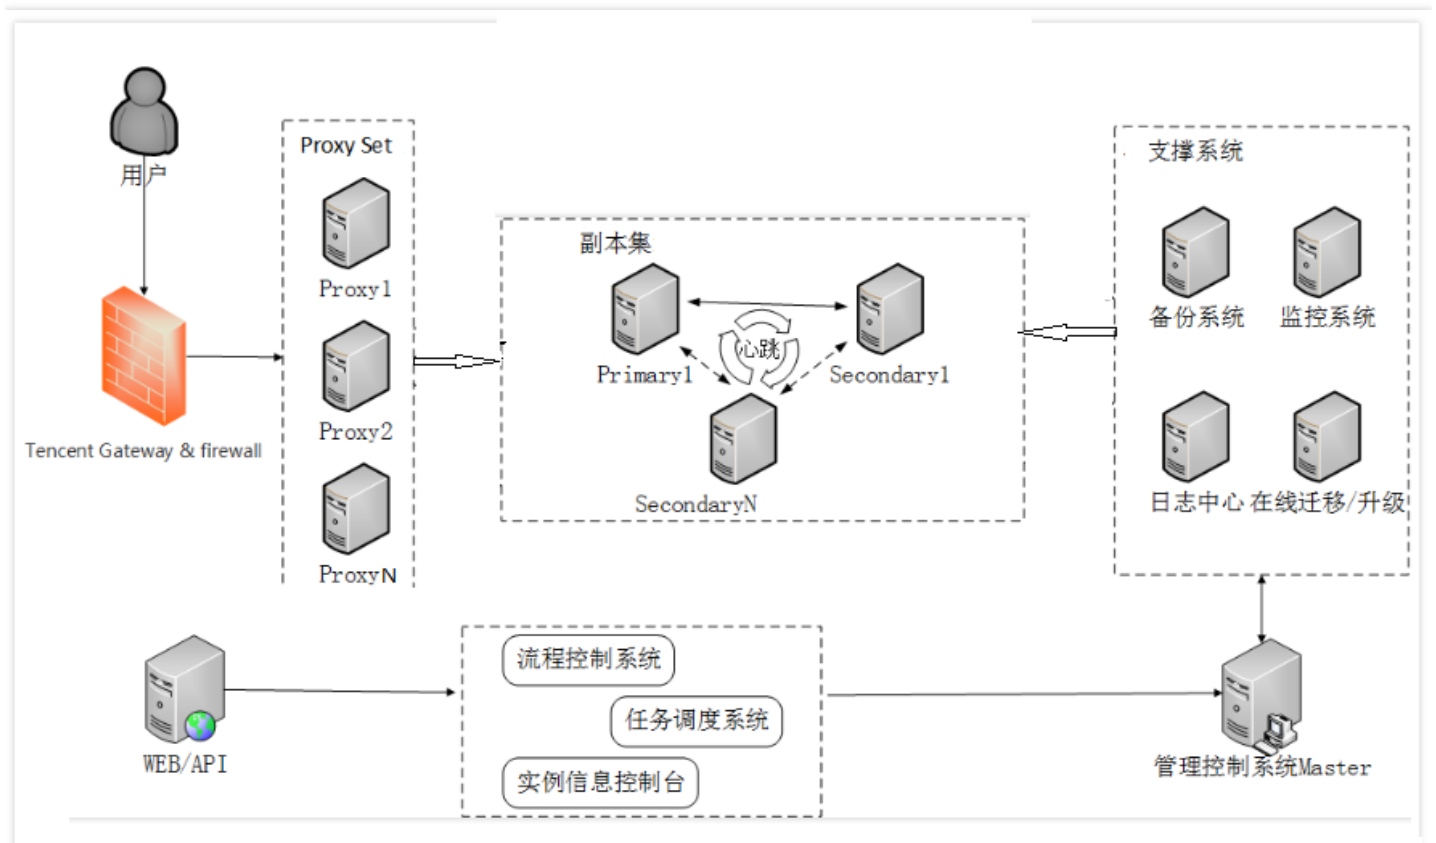
- (1) High flexibility and scalability. The database provides replica set and sharding instances. For replica set instances, nodes can be expanded vertically. For sharding instances, nodes can be expanded vertically and shards can be expanded horizontally to meet the expansion requirements of services.
- (2) Ultra-high performance. With super-large memory, new PCI-E SSD storage media and new generation storage engine, as well as the optimized Mongo kernel, it realizes a QPS of up to over 39,000 and can meet the requirements of massive and highly concurrent business scenarios.
- (3) High system availability. With a distributed architecture (at least 1 Primary and 1 Secondary), the system can automatically detect failures and switch services to ensure the high availability of the cluster. Meanwhile, the system performs a physical backup every 7 days and an incremental backup each day. You can roll back data to any time within 7 days using the file backups.
- (4) High data security. The isolation and high security of networks can be ensured via VPCs and security groups. Also, it provides the encrypted storage feature to ensure the high security of data at the storage layer, so the business is free of the data theft risk.
- (5) Data can be migrated. Support migration of instances with public network IPs, self-built migration of CVMs, direct connect migration and database migration. The business can freely move the database.
- (6) Read-only and remote disaster recovery. Provide read-only instances to meet the requirements of read/write separation in some business scenarios. Provide remote disaster recovery instances to quickly switch the service to the disaster recovery instance in case of a master instance failure to keep the service uninterrupted.

3. Product Architecture

A Replica set instance has a similar architecture with a sharding instance, except that there is only one shard in a replica set instance. The following describes the components in the system architecture.

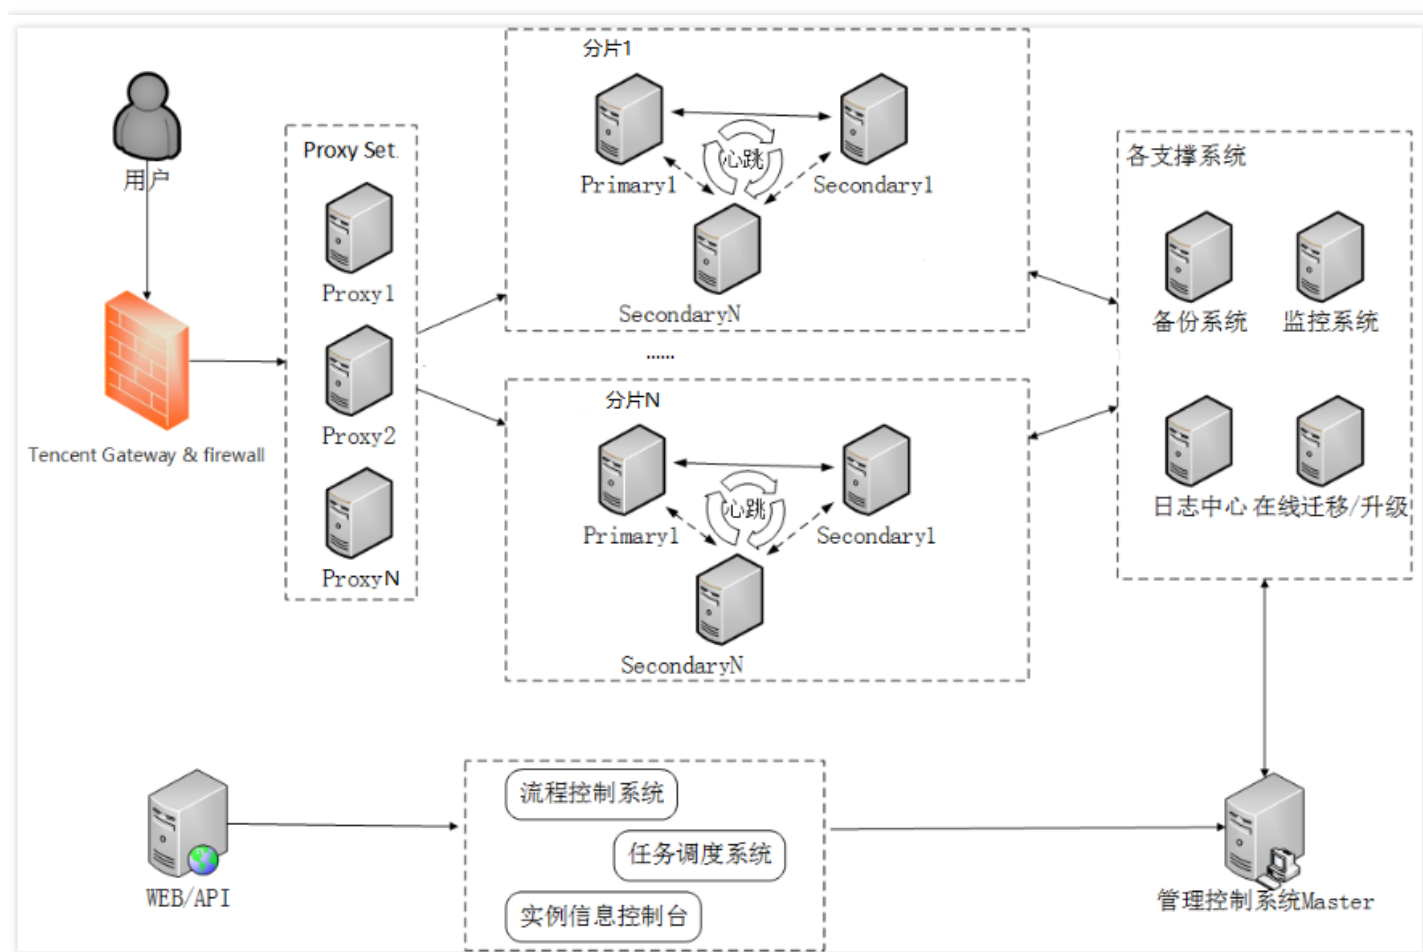
- Access module: ProxySet
ProxySet can provide consistent access for the replica set or the sharding cluster service. It is also partially responsible for reporting monitoring data and blocking sensitive operations.
- Management Control System: Master
It is the control center for the entire MongoDB cluster. It mainly manages the status and availability of each replica set or sharding cluster instance in the cluster, as well as providing functions such as migration, upgrade, backup, monitor, system deployment and so on.
- Monitor System: monitorsystem
It is mainly used to process monitoring data reported by each instance to analysis instance availability and reliability and push real-time alerts and e-mails to inform the status of instance under the username to the corresponding user.
- Backup System: backupcenter
It is used to store cold backup data of database MongoDB cluster.
- Log Center: logcenter
It is used to store the detailed access log for each replica set or sharding cluster instance and provide detailed log for detailed traceable issues.
- Task Scheduling System
It is responsible for the scheduling of task processes such as the creation, upgrade, and termination of instances in the entire system.

3.1 Replica set architecture



In the replica set architecture, the system automatically builds a "one master, two slaves" architecture. If the business expands, you can expand the computing or storage resources of each node.

3.2 Multipart architecture



In the sharding architecture, the system consists of multiple shards. If the business expands, the shards can be expanded indefinitely to support business scenarios of massive data.

Advantages

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

1. High performance

(1) With super-large memory, new PCI-E SSD storage media and new generation storage engine, it realizes a QPS of up to over 39,000.

(2) With a distributed architecture, it provides high system throughput to satisfy business scenarios of massive data.

2. High availability

(1) The distributed system architecture can prevent data loss.

(2) The system can realize automatic failure detection and recovery to ensure the high availability of instance.

(3) The system performs a full backup every 7 days and an incremental backup each day. You can roll back the business to any time within 7 days using the file backups.

(4) Provide read-only instances to meet the requirements in business scenarios of read/write separation. Provide remote disaster recovery instances to meet the strong demand for remote disaster recovery.

3. High security

(1) Provide VPC and security group features to ensure the security of business instances at the network layer.

(2) Provide the encrypted storage feature to ensure high security of business data.

(3) Provide a database-level account authorization and authentication mechanism to ensure the data security when the same database is used by different accounts of multiple businesses.

4. Ease of use

(1) You can easily realize the capacity expansion of node and shard expansion for instances via Tencent Cloud console.

(2) You can view the monitoring data of more than 20 metrics of instances via Tencent Cloud console, and can set up monitoring alarms to learn about the service usage details at any time.

Scenarios

Last updated : 2018-09-14 15:26:33

As a universal database, the stability, performance and scalability of TencentDB for MongoDB can satisfy most of the "no schema" scenarios. Typical application scenarios are as follows:

Game industry

MongoDB is especially suitable for game backend databases and can meet the fast-changing requirements of the game industry. MongoDB stores players' information, equipment, and credits in built-in documents to ensure convenient query and update. The "no schema" mode saves you the effort to change table structures and significantly shortens the version iteration cycle. MongoDB can also be used as a cache server to logically allocate hot data; its performance is comparable to other common cache servers while providing more query methods.

Mobile industry

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, as well as WeChat Moments and other information published by users.

IoT industry

Terminal devices in the field of IoT (Internet of Things), such as medical equipment, transport vehicle GPS, etc., can easily and continuously generate TBs of data. MongoDB can be used to store information of all the connected intelligent equipment, as well as log information reported by the equipment, and carry out multi-dimensional analysis on the information. You can build a distributed MongoDB sharding cluster to achieve storage without capacity limit and facilitate online capacity expansion so as to easily process massive IoT data.

Logistics industry

TencentDB for MongoDB can be used to store the order information. While the order status is continuously updated during delivery, MongoDB can store the order status in the embedded JSON format so that all order changes can be read in one query.

LVB industry

The LVB industry generates a large amount of gift information, user chat information and so on. TencentDB for MongoDB can be used to store user information, gift information, logs and so on. MongoDB also provides rich aggregate queries for business analysis.

Region Selection

Last updated : 2018-09-11 15:18:36

1. Regions

Tencent Cloud hosted data centers are distributed in different locations across the globe. Within China, they cover South China, East China and North China. The Hong Kong node covers Southeast Asia and the data center node in Toronto covers North America. We will gradually increase available nodes to cover more regions. All nodes consist of regions and availability zones.

Currently, the following available regions are offered by Tencent Cloud:

China: South China (Guangzhou), East China (Shanghai), North China (Beijing).

Outside China: Southeast Asia (Hong Kong) and North America (Toronto).

Notes:

- 1) Cloud products in the same region can communicate with each other through private network
- 2) Cloud products in different regions cannot communicate with each other through private network
- Tencent Cloud resources in different regions cannot access each other through private network
- 3) When purchasing cloud service, it is recommended to select the region that is closest to your customers to minimize access latency;
- 4) The name of regional availability zone is the most straightforward representation of the coverage range of a data center. To make it easier for clients to understand the name of regional availability zone, the "coverage range + city where the data center is located" structure is used when naming regions. The first half represents the coverage capability of the data center, and the second half indicates the city where the data center is located, or is near. Availability zone names are in the "city + number" structure.

Note:

- The "connection through private network" mentioned above refers to the interconnection among resources under the same account. Private resource networks of different accounts are completely isolated from each other
- Some Tencent Cloud resources may be unavailable in certain regions or availability zones. Before deploying your application, make sure you can create the required resources in that region or availability zone

2. Availability Zone

Availability zones refer to Tencent Cloud physical data centers in a certain region whose power and network equipments are independent from each other. They are designed to ensure that failures within an availability zone can be isolated (except for large-scale disaster or major power failure) without spreading to other zones in order to keep user business constantly up and running.

地域	可用区
华南地区（广州）	广州一区（售罄）
	广州二区
华东地区（上海）	上海一区
华北地区（北京）	北京一区
东南亚地区（香港）	香港一区
北美地区（多伦多）	多伦多一区

Availability zones in the same region can communicate with each other via private network, and network latency is shorter for products within the same availability zone.

Notes:

- If there is only one availability zone to choose in a region, it means the region has only one availability zone at the moment.
- You cannot change availability zone for cloud service resources already purchased.

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 : 2018-06-28 16:13:52

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.