TencentDB for MongoDB

Performance Data

Product Documentation

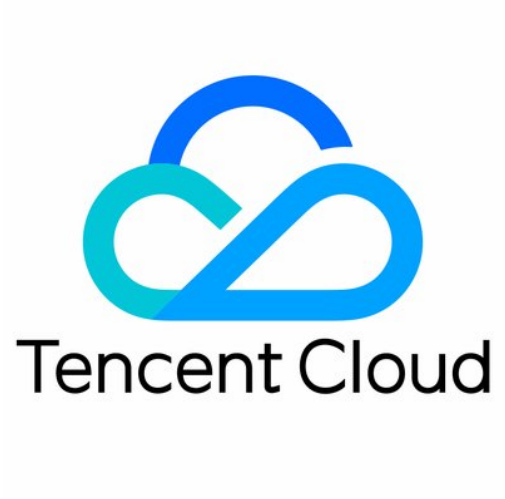

Performance Data

Last updated : 2021-03-17 19:58:41

This document describes how to perform standard performance testing on TencentDB for MongoDB instances. The resulting performance data is for your reference only.

Test Environment

- Test date: August 2020
- Client specification: the client is installed on an 8-core CVM instance with 32 GB memory. Tests show that if the MongoDB instance has a low specification, the CPU utilization of its replica instance can reach 100% under the request pressure of an 8-core CVM instance with 32 GB memory. In this case, one CVM instance achieves even better test results than multiple CVM instances can do. However, if the CPU utilization cannot reach 100% when you use only one CVM instance, you can try 4 CVM instances to share the concurrent threads.
- Test object: TencentDB for MongoDB v4.0 replica set instance

Testing Tool

Download YCSB.

Test Overview

Prepare about 10 GB data. Then, run 100 or 200 concurrent threads and use YCSB to test the throughput (ops/sec), RAL (read average latency in us), and WAL (write average latency in us) of instances with different specifications under the pressure of 50% read requests and 50% update requests, or 95% read requests and 5% update requests.

Latency

The average latency from CVM to MongoDB is 0.35 ms.
Minimum latency = 0.30 ms, maximum latency = 0.44 ms, average latency = 0.35 ms

Commands used in the test

1. Prepare data (about 10 GB)

```
nohup ./ycsb-0.15.0/bin/ycsb load mongodb -s -P workloads/workloada
-p mongodb.url=mongodb://mongouser:password@10.xx.xx.30:27017,10.xx.xx.28:27017,10.xx.xx.5:27017/
admin?w=0 -p table=test -threads 300 -p recordcount=10000000>loadlog.txt &
```

2. Test under 50% read requests and 50% update requests

```
nohup ./ycsb-0.15.0/bin/ycsb run mongodb -s -P workloads/workloada -p mongodb.url=mongodb://mongo
user: password @10.xx.xx.30:27017,10.xx.xx.28:27017,10.xx.xx.5:27017/admin?w=0 -p table=test -p r
ecordcount=10000000 -p readproportion=0.5 -p updateproportion=0.5 -p insertproportion=0 -p operat
ioncount=100000 -threads 100 >runlog.txt &
```

3. Test under 95% read requests and 5% update requests

```
nohup ./ycsb-0.15.0/bin/ycsb run mongodb -s -P workloads/workloada -p mongodb.url=mongodb://mongo
user: password @10.xx.xx.30:27017,10.xx.xx.28:27017,10.xx.xx.5:27017/admin?w=0 -p table=test -p r
ecordcount=10000000 -p readproportion=0.5 -p updateproportion=0.5 -p insertproportion=0 -p operat
ioncount=100000 -threads 100 >runlog.txt &
```

Note :

- Modify `-p operationcount=100000` according to the specific execution time to make sure that the execution time is longer than 20 minutes; otherwise, the resulting data will not be representative.
- `w` in `?w=0` represents write concern.
  - `w:1` (acknowledged write) requests acknowledgment that the write operation has propagated to the standalone mongod or the primary in a replica set. `w:1` is the default write concern for MongoDB.
  - `w:0` (unacknowledged write) requests no acknowledgment of the write operation. `w:0` will not report on the success of a write operation, but may return information about socket exceptions or networking errors when data is written to a closed socket or network is unavailable.
  - `w:>1` (valid in a replica set) requires acknowledgment from the primary and as many data-bearing secondaries as needed to meet the specified write concern.

Test Data

Read/update ratio (50:50)

MongoDB Specification	Threads	Throughput	RAL	WAL	CPU

		(ops/sec)	(us)	(us)	Utilization
2 cores, 4 GB memory	100	3,188	24,091	38,254	100%
2 cores, 4 GB memory	200	5,510	34,475	38,022	100%
4 cores, 8 GB memory	100	7,058	8,355	19,887	100%
4 cores, 8 GB memory	200	13,590	14,391	14,983	100%
6 cores, 16 GB memory	100	8,970	22,132	51	100%
6 cores, 16 GB memory	200	10,041	28,696	10,966	100%
12 cores, 32 GB memory	100	29,462	6,727	35	100%
12 cores, 32 GB memory	200	47,815	4,673	3,681	100%
24 cores, 64 GB memory	100	107,047	1,826	33	100%
24 cores, 64 GB memory	200	51,046	7,802	27	100%
24 cores, 128 GB memory	100	130,811	1,486	32	100%
24 cores, 128 GB memory	200	49,274	8,054	27	100%
32 cores, 240 GB memory	100	154,253	1,254	32	100%
32 cores, 240 GB memory	200	52,148	8,243	1,108	100%
48 cores, 512 GB memory	100	174,284	1,103	28	100%
48 cores, 512 GB memory	200	121,713	3,237	32	100%

Read/update ratio (95:5)

MongoDB Specification	Threads	Throughput (ops/sec)	RAL (us)	WAL (us)	CPU Utilization

2 cores, 4 GB memory	100	2,738	38,216	178	100%
2 cores, 4 GB memory	200	10,093	20,178	11,561	100%
4 cores, 8 GB memory	100	14,380	6,864	7,631	100%
4 cores, 8 GB memory	200	26,459	7,651	5,369	100%
6 cores, 16 GB memory	100	13,707	7,650	56	100%
6 cores, 16 GB memory	200	45,796	4,383	3,928	100%
12 cores, 32 GB memory	100	115,529	902	37	100%
12 cores, 32 GB memory	200	56,751	3,658	31	100%
24 cores, 64 GB memory	100	160,227	668	29	100%
24 cores, 64 GB memory	200	112,755	1,876	32	100%
24 cores, 128 GB memory	100	159,130	659	26	100%
24 cores, 128 GB memory	200	112,993	1,936	32	100%
32 cores, 240 GB memory	100	167,518	634	28	74%
32 cores, 240 GB memory	200	172,424	1,244	35	100%
48 cores, 512 GB memory	100	173,768	608	31	50%
48 cores, 512 GB memory	200	211,986	1,012	33	85%