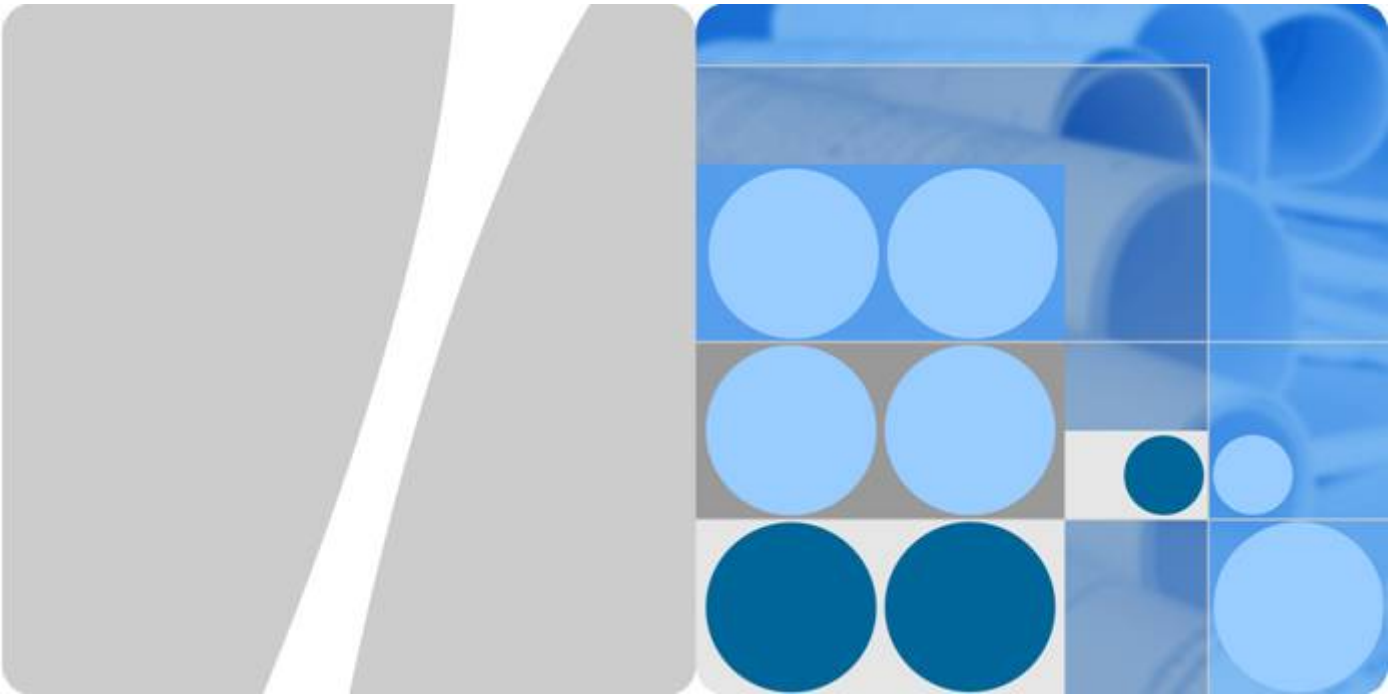

Document Database Service

Performance White Paper

Issue 01

Date 2020-11-20

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions


and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Contents

1 DDS Community Edition (General-Purpose)

1.1 Test Method

This section describes the performance test of DDS Community Edition 3.4, including the test environment, procedures, and results.

The following uses the cluster and replica set instances as an example.

Test Environment

- Region: CN North-Beijing1
- AZ: AZ1
- Elastic Cloud Server (ECS): s3.2xlarge.2 flavor with 8 vCPUs, 16 GB of memory, and CentOS 7.5 64 bit image.
- Each cluster instance includes two shard nodes.
- Specifications of the tested cluster and replica set instances: All specifications supported by the cluster and replica set instances are tested. For details, see **Table 1-1** and **Table 1-2**.

Table 1-1 Cluster instance class

Cluster ID	shard Class
Cluster 1	1 vCPUs | 4 GB
Cluster 2	2 vCPUs | 4 GB
Cluster 3	2 vCPUs | 8 GB
Cluster 4	4 vCPUs | 8 GB
Cluster 5	4 vCPUs | 16 GB
Cluster 6	8 vCPUs | 16 GB

Cluster ID	shard Class
Cluster 7	8 vCPUs | 32 GB

Table 1-2 Replica set instance class

Replica Set ID	Node Class
Replica set 1	1 vCPUs | 4 GB
Replica set 2	2 vCPUs | 4 GB
Replica set 3	2 vCPUs | 8 GB
Replica set 4	4 vCPUs | 8 GB
Replica set 5	4 vCPUs | 16 GB
Replica set 6	8 vCPUs | 16 GB
Replica set 7	8 vCPUs | 32 GB

Test Tool

YCSB is an open-source database performance test tool. In this test, YCSB 0.12.0 is used.

For details on how to use this tool, see **YCSB**.

Test Metrics

Operations per Second (OPS): number of operations executed per second by a database

Test Procedure

1. Configure the **workload** configuration file.

   Set the values of **readproportion**, **insertproportion**, and **updateproportion** in the **workload** file by referring to **Table 1-3**.

   Set the value of **recordcount** in the workload file based on the preset data volume listed in **Table 1-4**.

2. Use workload_s1 as an example. Run the following command to prepare test data:

   ./bin/ycsb load mongodb -s -P workloads/workload_s1 -p mongodb.url=mongodb://${userName}:${password}@${mongosIP}:${port}/ycsb?authSource=admin -threads ${threadNum} 1>workload_s1_load.result 2> workload_s1_load.log

3. Use workload_s1 as an example. Run the following command to test the performance:

   ./bin/ycsb run mongodb -s -P workloads/workload_s1 -p mongodb.url=mongodb://${userName}:${password}@${mongosIP}:$

{port}/ycsb?authSource=admin -threads ${threadNum} -p maxexecutiontime=1800 1>workload_s1_run.result 2> workload_s1_run.log

NOTE

- **${mongosIP}** indicates the private IP address of the mongos node in the DDS cluster instance.
- **${password}** indicates the administrator password of the DDS instance.
- **${threadNum}** indicates the number of concurrent threads for running the test. In this test, the number of concurrent threads is 128.

Testing Models

- Workload model

  Table 1-3 Service model

Service Model No.	Service Model
S1	100% insert
S2	90% update ,10% read
S3	65% read ,25% insert, 10% update
S4	90% read ,5% insert, 5% update
S5	50% update, 50% read
S6	100% read

- Number of concurrent threads: 128
- Document model

  Use the default configuration of YCSB: The size of each document is 1 KB, and the default index is **_id**.
- Data volume to be prepared

  In this test, prepare two types of data volume for each cluster instance.

  For details, see the following table.

  Table 1-4 Data volume to be prepared

Specifications	Low-Level Data Volume	High-Level Data Volume
1 vCPUs | 4 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
2 vCPUs | 4 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000

Specifications	Low-Level Data Volume	High-Level Data Volume
2 vCPUs | 8 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
4 vCPUs | 8 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
4 vCPUs | 16 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
8 vCPUs | 16 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
8 vCPUs | 32 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000

- Data consistency model

  Weak consistency: For the write concern setting of **{w: 1, j: false}**, an acknowledgment is returned after data is written to the disk on a single node. Data is persisted on disks in asynchronous mode with the default write concern setting.

1.2 Test Data

1.2.1 Cluster

NOTE

Each service model corresponds to a service model ID. For details, see **Service model**.

- Test OPS of the instance with different shard classes and the preset data of 10 million rows in different test service models. For details, see the content in bold in **Table 1-5**.

Table 1-5 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.							
shard Class	1 vCPUs | 4 GB	2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB

Service Model No.	s1	**9364**	**16437**	**16856**	**30619**	**37497**	**53742**	**55581**
	s6	**13715**	**14719**	**24688**	**38698**	**63828**	**80858**	**88973**
	s5	**7194**	**11203**	**12950**	**13517**	**25969**	**26918**	**47540**
	s4	**8503**	**12451**	**15306**	**17455**	**59859**	**86731**	**88534**
	s3	**9803**	**14242**	**17646**	**18568**	**35746**	**40218**	**63446**
	s2	**4827**	**8566**	**8689**	**8941**	**16938**	**17939**	**29237**

- Test OPS of the instance with different shard classes and the preset data of 100 million rows in different test service models. For details, see the content in bold in **Table 1-6**.

Table 1-6 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.								
shard Class		1U4GB	2U4GB	2U8GB	4U8GB	4U16GB	8U16GB	8U32GB
Service Model No.	s1	**5059**	**6967**	**9107**	**14808**	**16786**	**23125**	**29505**
	s6	**10120**	**14248**	**18216**	**18330**	**18504**	**18914**	**21089**
	s5	**6618**	**10974**	**11914**	**12066**	**12490**	**13117**	**13647**
	s4	**9067**	**12166**	**16322**	**16390**	**16446**	**16491**	**16810**
	s3	**6581**	**9232**	**11847**	**11856**	**15367**	**15679**	**15718**
	s2	**4276**	**6827**	**7697**	**9678**	**10271**	**10334**	**10762**

1.2.2 Replica Set

NOTE

Each service model corresponds to a service model ID. For details, see **Service model**.

- Test OPS of the instance with different classes and the preset data of 10 million rows in different test service models. For details, see the content in bold in **Table 1-7**.

Table 1-7 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.							
Node Class	1 vCPUs | 4 GB	2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB

Service Model No.	s1	**7830**	**11861**	**14094**	**26261**	**30532**	**50304**	**51178**
	s6	**4150**	**7446**	**7470**	**8050**	**20198**	**41831**	**79702**
	s5	**3703**	**5930**	**6666**	**7782**	**7741**	**8255**	**17033**
	s4	**4566**	**7857**	**8219**	**8337**	**11864**	**12302**	**75994**
	s3	**4888**	**8709**	**8800**	**8915**	**11550**	**12226**	**43370**
	s2	**2522**	**4710**	**4541**	**4769**	**5477**	**5496**	**10422**

- Test OPS of the instance with different classes and the preset data of 100 million rows in different test service models. For details, see the content in bold in **Table 1-8**.

Table 1-8 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.								
Node Class		1 vCPUs | 4 GB	2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB
Service Model No.	s1	**4413**	**6243**	**7945**	**12737**	**16840**	**23118**	**32035**
	s6	**3896**	**6855**	**7014**	**7039**	**9609**	**10011**	**10090**
	s5	**2997**	**5155**	**5395**	**5853**	**7313**	**7580**	**7606**
	s4	**3523**	**6152**	**6343**	**7033**	**8640**	**8684**	**9156**
	s3	**3758**	**6020**	**6766**	**7068**	**10682**	**10794**	**10855**
	s2	**1922**	**2040**	**3461**	**3514**	**3531**	**4147**	**5978**

2 DDS Community Edition (Enhanced)

2.1 Test Method

This section describes the performance test of enhanced instances of DDS Community Edition 3.4, including the test environment, procedures, and results.

The following uses the cluster and replica set instances as an example.

Test Environment

- Region: CN North-Beijing1
- AZ: AZ1
- Elastic Cloud Server (ECS): s3.2xlarge.2 flavor with 8 vCPUs, 16 GB of memory, and CentOS 7.5 64 bit image.
- Each cluster instance includes two shard nodes.
- For details about the DB instance class, see **Table 2-1** and **Table 2-2**.

Table 2-1 Cluster instance class

Cluster ID	shard Class
Cluster 1	1 vCPU | 4 GB
Cluster 2	2 vCPUs | 4 GB
Cluster 3	2 vCPUs | 8 GB
Cluster 4	4 vCPUs | 8 GB
Cluster 5	4 vCPUs | 16 GB
Cluster 6	8 vCPUs | 16 GB
Cluster 7	8 vCPUs | 32 GB
Cluster 8	16 vCPUs | 64 GB
Cluster 9	32 vCPUs | 128 GB

Cluster ID	shard Class
Cluster 10	60 vCPUs | 256 GB

Table 2-2 Replica set instance class

Cluster ID	shard Class
Replica set 1	1 vCPUs | 4 GB
Replica set 2	2 vCPUs | 4 GB
Replica set 3	2 vCPUs | 8 GB
Replica set 4	4 vCPUs | 8 GB
Replica set 5	4 vCPUs | 16 GB
Replica set 6	8 vCPUs | 16 GB
Replica set 7	8 vCPUs | 32 GB
Replica set 8	16 vCPUs | 64 GB
Replica set 9	32 vCPUs | 128 GB
Replica set 10	60 vCPUs | 256 GB

Test Tool

YCSB is an open-source database performance test tool. In this test, YCSB 0.12.0 is used.

For details on how to use this tool, see **YCSB**.

Test Metrics

Operations per Second (OPS): number of operations executed per second by a database

Test Procedure

1. Configure the **workload** configuration file.

   Set the values of **readproportion**, **insertproportion**, and **updateproportion** in the **workload** file by referring to **Table 2-3**.

   Set the value of **recordcount** in the **workload** file by referring to **Table 2-4**.

2. Use workload_s1 as an example. Run the following command to prepare test data:

   ./bin/ycsb load mongodb -s -P workloads/workload_s1 -p mongodb.url=mongodb://${userName}:${password}@${mongosIP}:${port}/ycsb?authSource=admin -threads ${threadNum} 1>workload_s1_load.result 2> workload_s1_load.log

3. Use workload_s1 as an example. Run the following command to test the performance:

   ./bin/ycsb run mongodb -s -P workloads/workload_s1 -p mongodb.url=mongodb://${userName}:${password}@${mongosIP}:${port}/ycsb?authSource=admin -threads ${threadNum} -p maxexecutiontime=1800 1>workload_s1_run.result 2> workload_s1_run.log

   NOTE

   - **${mongosIP}** indicates the private IP address of the mongos node in the DDS cluster instance.
   - **${password}** indicates the administrator password of the DDS instance.
   - **${threadNum}** indicates the number of concurrent threads for running the test. In this test, the number of concurrent threads is 128.

Testing Models

- Workload model

  Table 2-3 Service model

Service Model No.	Service Model
S1	100% insert
S2	90% update, 10% read
S3	65% read, 25% insert, 10% update
S4	90% read, 5% insert, 5% update
S5	50% update, 50% read
S6	100% read

- Number of concurrent threads: 128
- Document model

  Use the default configuration of YCSB: The size of each document is 1 KB, and the default index is **_id**.
- Data volume to be prepared

  In this test, prepare two types of data volume for each cluster instance.

  For details, see the following table.

  Table 2-4 Data volume to be prepared

Specifications	Low-Level Data Volume	High-Level Data Volume
1 vCPUs | 4 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000

Specifications	Low-Level Data Volume	High-Level Data Volume
2 vCPUs | 4 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
2 vCPUs | 8 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
4 vCPUs | 8 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
4 vCPUs | 16 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
8 vCPUs | 16 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
8 vCPUs | 32 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
16 vCPUs | 64 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
32 vCPUs | 128 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 10 GB Record counts: 10,000,000
60 vCPUs | 256 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 10 GB Record counts: 10,000,000

- Data consistency model

  Weak consistency: For the write concern setting of **{w: 1, j: false}**, an acknowledgment is returned after data is written to the disk on a single node. Data is persisted on disks in asynchronous mode with the default write concern setting.

2.2 Test Data

2.2.1 Cluster

NOTE

For details about the service model corresponding to the service model number, see **Table 2-3**.

- Test OPS of the instance with different shard classes and the preset data of 10 million rows in different test service models. For details, see the content in bold in **Table 2-5**.

Table 2-5 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.										
shard Class		1 vCPU | 4 GB	2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	60U 256 GB
Service Model No.	s1	**15610**	**17122**	**17559**	**31895**	**39131**	**56003**	**57943**	**65122**	**84123**	**108652**
	s6	**13708**	**15411**	**25719**	**40311**	**66488**	**84231**	**92680**	**96349**	**125002**	**160827**
	s5	**10098**	**11670**	**13498**	**14080**	**27081**	**28040**	**49532**	**54363**	**64981**	**78283**
	s4	**11935**	**12970**	**15944**	**18183**	**62361**	**90345**	**92236**	**102896**	**123178**	**160152**
	s3	**13006**	**14835**	**18382**	**19342**	**37235**	**41908**	**66090**	**73017**	**87097**	**113371**
	s2	**8625**	**8923**	**9053**	**9328**	**17644**	**18687**	**30455**	**38571**	**46216**	**50510**

- Test OPS of the instance with different shard classes and the preset data of 100 million rows in different test service models. For details, see the content in bold in **Table 2-6**.

Table 2-6 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.										
shard Class	1 vCPU | 4 GB	2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	60U 256 GB

Weak consistency: Test OPS of the database with 100 million rows of data.											
Service Model No.	s1	**4993**	**7258**	**9487**	**15425**	**17486**	**24092**	**30742**	**43523**	**56121**	**73231**
	s6	**9986**	**14847**	**18975**	**19096**	**19276**	**19703**	**21974**	**32967**	**42455**	**55215**
	s5	**6531**	**11467**	**12410**	**12569**	**13010**	**13664**	**14216**	**14578**	**15290**	**19098**
	s4	**8948**	**12679**	**17002**	**17074**	**17131**	**17259**	**17510**	**24247**	**31315**	**39481**
	s3	**6495**	**9619**	**12341**	**12354**	**16008**	**16333**	**16373**	**18927**	**24651**	**31472**
	s2	**4231**	**7112**	**8039**	**10029**	**10697**	**10829**	**11210**	**11297**	**14233**	**19091**

2.2.2 Replica Set

NOTE

For details about the service model corresponding to the service model number, see **Table 2-3**.

- Test OPS of the instance with different classes and the preset data of 10 million rows in different test service models. For details, see the content in bold in **Table 2-7**.

Table 2-7 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.											
DB Instance Class		1 vCPU | 4 GB	2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	60U 256 GB
Service Model No.	s1	**12165**	**12355**	**14598**	**27287**	**31832**	**52401**	**53310**	**59304**	**76736**	**98928**
	s6	**7625**	**7756**	**7781**	**8385**	**21040**	**43513**	**83024**	**97673**	**125179**	**163671**
	s5	**6147**	**6203**	**6944**	**8107**	**8074**	**8596**	**17748**	**46933**	**61012**	**79300**

Weak consistency: Test OPS of the database with 10 million rows of data.											
	s4	**8236**	**8248**	**8559**	**8688**	**12269**	**12815**	**79160**	**98156**	**127601**	**165655**
	s3	**9007**	**9072**	**9167**	**9287**	**12032**	**12735**	**45178**	**72371**	**94023**	**122304**
	s2	**4900**	**4906**	**4956**	**4968**	**5701**	**5725**	**10857**	**26776**	**34808**	**45250**

- Test OPS of the instance with different classes and the preset data of 100 million rows in different test service models. For details, see the content in bold in **Table 2-8**.

Table 2-8 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.											
DB Instance Class		1 vCPU | 4 GB	2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	60U 256 GB
Service Model No.	s1	**4358**	**6504**	**8281**	**13286**	**17542**	**24152**	**33424**	**36815**	**47840**	**62165**
	s6	**3845**	**7141**	**7306**	**7307**	**10037**	**10468**	**10514**	**10520**	**13676**	**17778**
	s5	**2957**	**5370**	**5620**	**6097**	**7617**	**7896**	**7924**	**8001**	**10287**	**13519**
	s4	**3674**	**6124**	**6981**	**7302**	**9001**	**9046**	**9538**	**9608**	**12487**	**16235**
	s3	**3709**	**6271**	**7048**	**7363**	**11127**	**11244**	**11308**	**11344**	**14746**	**19171**
	s2	**1897**	**2125**	**3605**	**3660**	**3679**	**4320**	**6227**	**6490**	**8437**	**10968**

3 DDS Community Edition (Enhanced II)

3.1 Test Method

This section describes the performance test of DDS Enhanced II of Community Edition 3.4 and 4.0, including the test environment, procedures, and results.

The following uses the cluster and replica set instances as an example.

Test Environment

- Region: CN North-Beijing4
- AZ: AZ2
- Elastic Cloud Server (ECS): c3.2xlarge.2 flavor with 8 vCPUs, 16 GB of memory, and CentOS 7.6 (x86, 64 bit).
- DB engine versions of the cluster and replica set instances to be tested: 3.4 and 4.0.
- Each cluster instance includes two shard nodes.
- Specifications of the cluster and replica set Enhanced II instances to be tested: All the specifications supported by the instances are covered. For details, see **Table 3-1** and **Table 3-2**.

Table 3-1 Cluster instance specifications

ID	DB Version	shard Class
Cluster 1	3.4	2 vCPUs | 4 GB
Cluster 2	3.4	2 vCPUs | 8 GB
Cluster 3	3.4	4 vCPUs | 8 GB
Cluster 4	3.4	4 vCPUs | 16 GB
Cluster 5	3.4	8 vCPUs | 16 GB
Cluster 6	3.4	8 vCPUs | 32 GB

ID	DB Version	shard Class
Cluster 7	3.4	16 vCPUs | 32 GB
Cluster 8	3.4	16 vCPUs | 64 GB
Cluster 9	3.4	32 vCPUs | 128 GB
Cluster 10	3.4	64 vCPUs | 256 GB
Cluster 11	4.0	2 vCPUs | 4 GB
Cluster 12	4.0	2 vCPUs | 8 GB
Cluster 13	4.0	4 vCPUs | 8 GB
Cluster 14	4.0	4 vCPUs | 16 GB
Cluster 15	4.0	8 vCPUs | 16 GB
Cluster 16	4.0	8 vCPUs | 32 GB
Cluster 17	4.0	16 vCPUs | 32 GB
Cluster 18	4.0	16 vCPUs | 64 GB
Cluster 19	4.0	32 vCPUs | 128 GB
Cluster 20	4.0	64 vCPUs | 256 GB

Table 3-2 Replica set specifications

ID	DB Version	DB Instance Class
Replica set 1	3.4	2 vCPUs | 4 GB
Replica set 2	3.4	2 vCPUs | 8 GB
Replica set 3	3.4	4 vCPUs | 8 GB
Replica set 4	3.4	4 vCPUs | 16 GB
Replica set 5	3.4	8 vCPUs | 16 GB
Replica set 6	3.4	8 vCPUs | 32 GB
Replica set 7	3.4	16 vCPUs | 32 GB
Replica set 8	3.4	16 vCPUs | 64 GB
Replica set 9	3.4	32 vCPUs | 128 GB
Replica set 10	3.4	64 vCPUs | 256 GB
Replica set 11	4.0	2 vCPUs | 4 GB
Replica set 12	4.0	2 vCPUs | 8 GB

ID	DB Version	DB Instance Class
Replica set 13	4.0	4 vCPUs | 8 GB
Replica set 14	4.0	4 vCPUs | 16 GB
Replica set 15	4.0	8 vCPUs | 16 GB
Replica set 16	4.0	8 vCPUs | 32 GB
Replica set 17	4.0	16 vCPUs | 32 GB
Replica set 18	4.0	16 vCPUs | 64 GB
Replica set 19	4.0	32 vCPUs | 128 GB
Replica set 20	4.0	64 vCPUs | 256 GB

Test Tool

YCSB is an open-source database performance test tool. In this test, YCSB **0.12.0** is used.

For details on how to use this tool, see **YCSB**.

Test Metrics

Operations per Second (OPS): number of operations executed per second by a database

Test Procedure

1. Configure the **workload** configuration file.

   Set the values of **readproportion**, **insertproportion**, and **updateproportion** in the **workload** file by referring to **Table 3-3**.

   Set the value of **recordcount** in the **workload** file by referring to **Table 3-4**.

   Example: Configure the **workload_s1** file.

   - recordcount = 100000000
   - operationcount = 100000000
   - insertproportion = 1
   - readproportion = 0
   - updateproportion = 0
   - scanproportion = 0

   NOTE

   The values of **recordcount** and **operationcount** are the same.

   The sum of the values of **insertproportion**, **readproportion**, **updateproportion**, and **scanproportion** is **1**.

2. Run the following command to create a test database table on the DB instance, for example, **usertable**:

   ./bin/mongo --host $mongosIP --port $port << END_CMD

use admin

db.auth('rwuser', '$pwd')

use ycsb

sh.enableSharding("ycsb")

sh.shardCollection("ycsb.usertable",{_id:"hashed"},false, {numInitialChunks:200})

END_CMD

3. Use workload_s1 and cluster instances as an example. Run the following command to prepare data:

   ./bin/ycsb load mongodb -s -P workloads/workload_s1 -p mongodb.url=mongodb://rwuser:${password}@${mongosIP}:8635,${mongosIP}:8635/ycsb?authSource=admin -threads ${threadNum} 1>workload_s1_load.result 2> workload_s1_load.log

4. Use workload_s1 and cluster instances as an example. Run the following command to test the performance:

   ./bin/ycsb run mongodb -s -P workloads/workload_s1 -p mongodb.url=mongodb://rwuser:${password}@${mongosIP}:8635,${mongosIP}:8635/ycsb?authSource=admin -threads ${threadNum} -p maxexecutiontime=1800 1>workload_s1_run.result 2> workload_s1_run.log

   NOTE

   - **${mongosIP}** indicates the private IP address of the mongos node in the DDS cluster instance.
   - **${password}** indicates the administrator password of the DDS instance.
   - **${threadNum}** indicates the number of concurrent threads for running the test. In this test, the number of concurrent threads is 128.

Testing Models

- Workload model

  Table 3-3 Service model

Service Model No.	Service Model
S1	100% insert
S2	90% update ,10% read
S3	65% read ,25% insert, 10% update
S4	90% read ,5% insert, 5% update
S5	50% update, 50% read
S6	100% read

- Number of concurrent threads: 128
- Document model

Use the default configuration of YCSB: The size of each document is 1 KB, and the default index is **_id**.

- Data volume to be prepared

  In this test, prepare two types of data volume for each cluster instance.

  For details, see the following table.

  Table 3-4 Data volume to be prepared

Specifications	Low-Level Data Volume	High-Level Data Volume
2 vCPUs | 4 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
2 vCPUs | 8 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
4 vCPUs | 8 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
4 vCPUs | 16 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
8 vCPUs | 16 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
8 vCPUs | 32 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
16 vCPUs | 32 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
16 vCPUs | 64 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
32 vCPUs | 128 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000
64 vCPUs | 256 GB	Storage space: 10 GB Record counts: 10,000,000	Storage space: 100 GB Record counts: 100,000,000

- Data consistency model

Weak consistency: For the write concern setting of **{w: 1, j: false}**, an acknowledgment is returned after data is written to the disk on a single node. Data is persisted on disks in asynchronous mode with the default write concern setting.

3.2 Test Data

3.2.1 Cluster

NOTE

For details about the service model corresponding to the service model number, see **Table 2-3**.

- Test OPS of DB instances of version 3.4 with different shard classes and the preset data of 10 million rows in different test service models. For details, see the content in bold in **Table 3-5**.

Table 3-5 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.									
shard Class		2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	64U 256 GB
Service Model No.	s1	**20546**	**21070**	**38274**	**46871**	**67178**	**69476**	**78146**	**100947**	**130382**
	s6	**18399**	**30860**	**48373**	**79785**	**101073**	**111216**	**115603**	**150002**	**192911**
	s5	**14004**	**16188**	**16896**	**32461**	**33648**	**59425**	**65235**	**77951**	**93939**
	s4	**15564**	**19132**	**21819**	**74824**	**108414**	**110668**	**123473**	**147812**	**192024**
	s3	**17802**	**22058**	**23210**	**44682**	**50272**	**79308**	**87620**	**104516**	**136045**
	s2	**10707**	**10861**	**11176**	**21172**	**22424**	**36546**	**46229**	**55456**	**60612**

- Test OPS of DB instances of version 3.4 with different shard classes and the preset data of 100 million rows in different test service models. For details, see the content in bold in **Table 3-6**.

Table 3-6 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.										
shard Class		2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	64U 256 GB
Service Model No.	s1	**87009**	**11384**	**18510**	**20982**	**28906**	**36881**	**52220**	**67345**	**87877**
	s6	**17810**	**22770**	**22912**	**23130**	**23642**	**26361**	**39560**	**50946**	**66258**
	s5	**13717**	**14892**	**15082**	**15612**	**16396**	**17059**	**17493**	**18348**	**22917**
	s4	**15208**	**20402**	**20488**	**20557**	**20614**	**21012**	**29034**	**37578**	**47367**
	s3	**11540**	**14809**	**14820**	**19209**	**19599**	**19647**	**22123**	**29329**	**37631**
	s2	**8534**	**9621**	**12098**	**12839**	**12917**	**13452**	**13553**	**17082**	**22841**

- Test OPS of DB instances of version 4.0 with different shard classes and the preset data of 10 million rows in different test service models. For details, see the content in bold in **Table 3-7**.

Table 3-7 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.										
shard Class		2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	64U 256 GB
Service Model No.	s1	**19930**	**20438**	**37126**	**45465**	**65163**	**67392**	**75802**	**97919**	**126471**
	s6	**18031**	**30243**	**47406**	**78189**	**99052**	**108992**	**113291**	**147002**	**189053**
	s5	**13654**	**15783**	**16474**	**31649**	**32807**	**57939**	**63604**	**76002**	**91591**
	s4	**15268**	**18768**	**21404**	**73402**	**106354**	**108565**	**121127**	**145004**	**188376**

Weak consistency: Test OPS of the database with 10 million rows of data.										
	s3	**17482**	**21661**	**22792**	**43878**	**49367**	**77880**	**86043**	**102635**	**133596**
	s2	**10471**	**10622**	**10930**	**20706**	**21931**	**35742**	**45212**	**54236**	**59279**

- Test OPS of DB instances of version 4.0 with different shard classes and the preset data of 100 million rows in different test service models. For details, see the content in bold in **Table 3-8**.

Table 3-8 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.										
shard Class		2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	64U 256 GB
Service Model No.	s1	**8448**	**11042**	**17955**	**20353**	**28039**	**35775**	**50653**	**65325**	**85241**
	s2	**17454**	**22315**	**22454**	**22667**	**23169**	**25834**	**38769**	**49927**	**64933**
	s3	**13374**	**14520**	**14705**	**15222**	**15986**	**16633**	**17056**	**17889**	**22344**
	s4	**14919**	**20014**	**20099**	**20166**	**20222**	**20613**	**28482**	**36864**	**46467**
	s5	**11332**	**14542**	**14553**	**18863**	**19246**	**19293**	**21725**	**28801**	**36954**
	s6	**8346**	**9409**	**11832**	**12557**	**12633**	**13156**	**13255**	**16706**	**22338**

3.2.2 Replica Set

NOTE

For details about the service model corresponding to the service model number, see **Table 2-3**.

- Test OPS of DB instances of version 3.4 with different classes and the preset data of 10 million rows in different test service models. For details, see the content in bold in **Table 3-9**.

Table 3-9 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.										
DB Instance Class		2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	64U 256 GB
Service Model No.	s1	**14826**	**17617**	**32826**	**38165**	**62880**	**63972**	**71114**	**92089**	**118709**
	s6	**9307**	**9337**	**10062**	**25248**	**52289**	**99628**	**117312**	**150876**	**193819**
	s5	**7412**	**8332**	**9728**	**9676**	**10319**	**21291**	**56314**	**73388**	**95160**
	s4	**9821**	**10274**	**10421**	**14830**	**15378**	**94992**	**117787**	**153121**	**198786**
	s3	**10886**	**11000**	**11144**	**14438**	**15282**	**54213**	**86845**	**112827**	**146877**
	s2	**5887**	**5676**	**5961**	**6846**	**6870**	**13028**	**32131**	**41769**	**54300**

- Test OPS of DB instances of version 3.4 with different classes and the preset data of 100 million rows in different test service models. For details, see the content in bold in **Table 3-10**.

Table 3-10 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.										
DB Instance Class		2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	64U 256 GB
Service Model No.	s1	**7804**	**9931**	**15921**	**21050**	**28897**	**40044**	**44178**	**57408**	**74598**
	s6	**8569**	**8767**	**8799**	**12011**	**12514**	**12612**	**12624**	**16417**	**21332**
	s5	**6444**	**6744**	**7316**	**9141**	**9475**	**9508**	**9601**	**12481**	**16203**
	s4	**7231**	**8012**	**8835**	**10800**	**10855**	**11445**	**11529**	**14984**	**19482**

Weak consistency: Test OPS of the database with 100 million rows of data.										
	s3	**7525**	**8457**	**8835**	**13352**	**13492**	**13569**	**13612**	**17695**	**23005**
	s2	**2550**	**4326**	**4392**	**4414**	**5184**	**7472**	**7788**	**10124**	**13161**

- Test OPS of DB instances of version 4.0 with different classes and the preset data of 10 million rows in different test service models. For details, see the content in bold in **Table 3-11**.

Table 3-11 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.										
DB Instance Class		2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	64U 256 GB
Service Model No.	s1	**14381**	**17088**	**31841**	**37020**	**60994**	**62053**	**68981**	**89326**	**115148**
	s2	**9121**	**9150**	**9861**	**24743**	**51243**	**97635**	**114966**	**147858**	**189943**
	s3	**7227**	**8124**	**9485**	**9434**	**10061**	**20759**	**54906**	**71553**	**92781**
	s4	**9634**	**10079**	**10223**	**14548**	**15086**	**93187**	**115549**	**150212**	**195009**
	s5	**10690**	**10802**	**10943**	**14178**	**15007**	**53237**	**85282**	**110796**	**144233**
	s6	**5757**	**5551**	**5830**	**6695**	**6719**	**12741**	**31424**	**40850**	**53105**

- Test OPS of DB instances of version 4.0 with different classes and the preset data of 100 million rows in different test service models. For details, see the content in bold in **Table 3-12**.

Table 3-12 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.										
DB Instance Class		2 vCPUs | 4 GB	2 vCPUs | 8 GB	4 vCPUs | 8 GB	4 vCPUs | 16 GB	8 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32U 128 GB	64U 256 GB
Service Model No.	s1	**7570**	**9633**	**15443**	**20419**	**28030**	**38843**	**42853**	**55686**	**72360**
	s2	**8398**	**8592**	**8623**	**11771**	**12264**	**12360**	**12372**	**16089**	**20905**
	s3	**6283**	**6575**	**7133**	**8912**	**9238**	**9270**	**9361**	**12169**	**15798**
	s4	**7094**	**7860**	**8667**	**10595**	**10649**	**11228**	**11310**	**14699**	**19112**
	s5	**7390**	**8305**	**8676**	**13112**	**13249**	**13325**	**13367**	**17376**	**22591**
	s6	**2494**	**4231**	**4295**	**4317**	**5070**	**7308**	**7617**	**9901**	**12871**