

Document Database Service

Performance White Paper

Issue 01
Date 2022-08-30



Copyright © Huawei Technologies Co., Ltd. 2022. All rights reserved.

No part of this document may be reproduced or transmitted in any form or by any means without prior written consent of Huawei Technologies Co., Ltd.

Trademarks and Permissions



HUAWEI 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
1.1 Test Method.....	1
1.2 Test Data.....	5
1.2.1 Cluster.....	5
1.2.2 Replica Set.....	7
2 DDS Community Edition (Enhanced II)	11
2.1 Test Method.....	11
2.2 Test Data.....	16
2.2.1 Cluster.....	16
2.2.2 Replica Set.....	20

1 DDS Community Edition (General-Purpose)

1.1 Test Method

This section describes the performance test of DDS 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

- 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	DB Version	shard Class	Storage Space
Cluster 1	3.4	1 vCPU 4 GB	600 GB
Cluster 2	3.4	2 vCPUs 4 GB	600 GB
Cluster 3	3.4	2 vCPUs 8 GB	600 GB
Cluster 4	3.4	4 vCPUs 8 GB	600 GB
Cluster 5	3.4	4 vCPUs 16 GB	600GB
Cluster 6	3.4	8 vCPUs 16 GB	600 GB
Cluster 7	3.4	8 vCPUs 32 GB	600 GB

Cluster ID	DB Version	shard Class	Storage Space
Cluster 8	4.0	1 vCPU 4 GB	600 GB
Cluster 9	4.0	2 vCPUs 4 GB	600 GB
Cluster 10	4.0	2 vCPUs 8 GB	600 GB
Cluster 11	4.0	4 vCPUs 8 GB	600 GB
Cluster 12	4.0	4 vCPUs 16 GB	600 GB
Cluster 13	4.0	8 vCPUs 16 GB	600 GB
Cluster 14	4.0	8 vCPUs 32 GB	600 GB

Table 1-2 Replica set instance class

Replica Set ID	DB Version	Node Class	Storage Space
Replica set 1	3.4	1 vCPU 4 GB	600 GB
Replica set 2	3.4	2 vCPUs 4 GB	600 GB
Replica set 3	3.4	2 vCPUs 8 GB	600 GB
Replica set 4	3.4	4 vCPUs 8 GB	600 GB
Replica set 5	3.4	4 vCPUs 16 GB	600 GB
Replica set 6	3.4	8 vCPUs 16 GB	600 GB
Replica set 7	3.4	8 vCPUs 32 GB	600 GB
Replica set 8	4.0	1 vCPU 4 GB	600 GB
Replica set 9	4.0	2 vCPUs 4 GB	600 GB
Replica set 10	4.0	2 vCPUs 8 GB	600 GB
Replica set 11	4.0	4 vCPUs 8 GB	600 GB
Replica set 12	4.0	4 vCPUs 16 GB	600 GB
Replica set 13	4.0	8 vCPUs 16 GB	600 GB
Replica set 14	4.0	8 vCPUs 32 GB	600 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](#).

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

```
./bin/yycsb load mongodb -s -P workloads/workload_s1 -p  
mongodb.url="mongodb://{userName}:{password}@{mongosIP}:{  
{port}/yycsb?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/yycsb run mongodb -s -P workloads/workload_s1 -p  
mongodb.url="mongodb://{userName}:{password}@{mongosIP}:{  
{port}/yycsb?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
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

Specifications	Low-Level Data Volume	High-Level Data Volume
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](#).
- Operations per Second (OPS): number of operations executed per second by a database
- In different service models, 10 million rows of data are preset for DB instances 3.4 with different shard specifications. For details about the test OPS, see the content in bold in [Table 1-5](#).

Table 1-5 Weak consistency: Test OPS of the database with 10 million rows of data.

shard Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
1 vCPU 4 GB	9364	13715	7194	8503	9803	4827
2 vCPUs 4 GB	16437	14719	11203	12451	14242	8566
2 vCPUs 8 GB	16856	24688	12950	15306	17646	8689
4 vCPUs 8 GB	30619	38698	13517	17455	18568	8941
4 vCPUs 16 GB	37497	63828	25969	59859	35746	16938
8 vCPUs 16 GB	53742	80858	26918	86731	40218	17939
8 vCPUs 32 GB	55581	88973	47540	88534	63446	29237

- In different service models, 100 million rows of data are preset for DB instances 3.4 with different shard specifications. For details about the test OPS, see the content in bold in [Table 1-6](#).

Table 1-6 Weak consistency: Test OPS of the database with 100 million rows of data.

shard Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
1 vCPU 4 GB	5059	10120	6618	9067	6581	4276
2 vCPUs 4 GB	6967	14248	10974	12166	9232	6827
2 vCPUs 8 GB	9107	18216	11914	16322	11847	7697
4 vCPUs 8 GB	14808	18330	12066	16390	11856	9678
4 vCPUs 16 GB	16786	18504	12490	16446	15367	10271
8 vCPUs 16 GB	23125	18914	13117	16491	15679	10334
8 vCPUs 32 GB	29505	21089	13647	16810	15718	10762

- In different service models, 10 million rows of data are preset for DB instances 4.0 with different shard specifications. For details about the test OPS, see the content in bold in [Table 1-7](#).

Table 1-7 Weak consistency: Test OPS of the database with 10 million rows of data.

shard Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
1 vCPU 4 GB	9299	13520	7194	8464	9725	4898
2 vCPUs 4 GB	16160	14493	11083	12293	14031	8525
2 vCPUs 8 GB	16566	24163	12778	15063	17333	8644
4 vCPUs 8 GB	29916	37753	13327	17147	18227	8889

4 vCPUs 16 GB	36588	62129	25406	58279	34890	16646
8 vCPUs 16 GB	52346	78648	26326	84345	39227	17617
8 vCPUs 32 GB	54130	86520	46330	86094	61759	28576

- In different service models, 100 million rows of data are preset for DB instances 4.0 with different shard specifications. For details about the test OPS, see the content in bold in [Table 1-8](#).

Table 1-8 Weak consistency: Test OPS of the database with 100 million rows of data.

shard Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
1 vCPU 4 GB	5123	10032	6635	9011	6600	4364
2 vCPUs 4 GB	6974	14037	10861	12017	9171	6838
2 vCPUs 8 GB	9050	17886	11773	16048	11708	7682
4 vCPUs 8 GB	14580	17996	11920	16114	11716	9604
4 vCPUs 16 GB	16498	18165	12331	16169	15122	10179
8 vCPUs 16 GB	22647	18563	12939	16212	15425	10240
8 vCPUs 32 GB	28836	20672	13454	16522	15462	10655

1.2.2 Replica Set

NOTE

- Each service model corresponds to a service model ID. For details, see [Service model](#).
- Operations per Second (OPS): number of operations executed per second by a database
- In different service models, 10 million rows of data are preset for DB instances 3.4 with different specifications. For details about the test OPS, see the content in bold in [Table 1-9](#).

Table 1-9 Weak consistency: Test OPS of the database with 10 million rows of data.

DB Instance Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
1 vCPU 4 GB	7830	4150	3703	4566	4888	2522
2 vCPUs 4 GB	11861	7446	5930	7857	8709	4710
2 vCPUs 8 GB	14094	7470	6666	8219	8800	4541
4 vCPUs 8 GB	26261	8050	7782	8337	8915	4769
4 vCPUs 16 GB	30532	20198	7741	11864	11550	5477
8 vCPUs 16 GB	50304	41831	8255	12302	12226	5496
8 vCPUs 32 GB	51178	79702	17033	75994	43370	10422

- In different service models, 100 million rows of data are preset for DB instances 3.4 with different specifications. For details about the test OPS, see the content in bold in [Table 1-10](#).

Table 1-10 Weak consistency: Test OPS of the database with 100 million rows of data.

DB Instance Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
1 vCPU 4 GB	4413	3896	2997	3523	3758	1922
2 vCPUs 4 GB	6243	6855	5155	6152	6020	2040
2 vCPUs 8 GB	7945	7014	5395	6343	6766	3461
4 vCPUs 8 GB	12737	7039	5853	7033	7068	3514
4 vCPUs 16 GB	16840	9609	7313	8640	10682	3531

8 vCPUs 16 GB	23118	10011	7580	8684	10794	4147
8 vCPUs 32 GB	32035	10090	7606	9156	10855	5978

- In different service models, 10 million rows of data are preset for DB instances 4.0 with different specifications. For details about the test OPS, see the content in bold in [Table 1-11](#).

Table 1-11 Weak consistency: Test OPS of the database with 10 million rows of data.

DB Instance Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
1 vCPU 4 GB	7811	4242	3808	4645	4957	2662
2 vCPUs 4 GB	11721	7439	5968	7837	8664	4785
2 vCPUs 8 GB	13887	7462	6682	8188	8752	4621
4 vCPUs 8 GB	25689	8025	7765	8303	8864	4842
4 vCPUs 16 GB	29832	19808	7725	11724	11420	5529
8 vCPUs 16 GB	49011	40792	8223	12149	12075	5547
8 vCPUs 32 GB	49859	77527	16738	73930	42285	10325

- In different service models, 100 million rows of data are preset for DB instances 4.0 with different specifications. For details about the test OPS, see the content in bold in [Table 1-12](#).

Table 1-12 Weak consistency: Test OPS of the database with 100 million rows of data.

DB Instance Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
1 vCPU 4 GB	4497	3995	3123	3633	3861	2080

2 vCPUs 4 GB	6272	6865	5216	6183	6055	2195
2 vCPUs 8 GB	7923	7020	5449	6369	6779	3573
4 vCPUs 8 GB	12571	7044	5893	7038	7072	3625
4 vCPUs 16 GB	16551	9537	7310	8597	10578	3641
8 vCPUs 16 GB	22640	9927	7569	8639	10686	4239
8 vCPUs 32 GB	31290	10003	7594	9097	10745	6015

2 DDS Community Edition (Enhanced II)

2.1 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

- 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 2-1](#) and [Table 2-2](#).

Table 2-1 Cluster instance specifications

ID	DB Version	shard Class	Storage Space
Cluster 1	3.4	2 vCPUs 4 GB	600 GB
Cluster 2	3.4	2 vCPUs 8 GB	600 GB
Cluster 3	3.4	4 vCPUs 8 GB	600 GB
Cluster 4	3.4	4 vCPUs 16 GB	600 GB
Cluster 5	3.4	8 vCPUs 16 GB	600 GB
Cluster 6	3.4	8 vCPUs 32 GB	600 GB
Cluster 7	3.4	16 vCPUs 32 GB	600 GB

ID	DB Version	shard Class	Storage Space
Cluster 8	3.4	16 vCPUs 64 GB	600 GB
Cluster 9	3.4	32 vCPUs 128 GB	600 GB
Cluster 10	3.4	64 vCPUs 256 GB	600 GB
Cluster 11	4.0	2 vCPUs 4 GB	600 GB
Cluster 12	4.0	2 vCPUs 8 GB	600 GB
Cluster 13	4.0	4 vCPUs 8 GB	600 GB
Cluster 14	4.0	4 vCPUs 16 GB	600 GB
Cluster 15	4.0	8 vCPUs 16 GB	600 GB
Cluster 16	4.0	8 vCPUs 32 GB	600 GB
Cluster 17	4.0	16 vCPUs 32 GB	600 GB
Cluster 18	4.0	16 vCPUs 64 GB	600 GB
Cluster 19	4.0	32 vCPUs 128 GB	600 GB
Cluster 20	4.0	64 vCPUs 256 GB	600 GB

Table 2-2 Replica set specifications

ID	DB Version	DB Instance Class	Storage Space
Replica set 1	3.4	2 vCPUs 4 GB	600 GB
Replica set 2	3.4	2 vCPUs 8 GB	600 GB
Replica set 3	3.4	4 vCPUs 8 GB	600 GB
Replica set 4	3.4	4 vCPUs 16 GB	600 GB
Replica set 5	3.4	8 vCPUs 16 GB	600 GB
Replica set 6	3.4	8 vCPUs 32 GB	600 GB
Replica set 7	3.4	16 vCPUs 32 GB	600 GB
Replica set 8	3.4	16 vCPUs 64 GB	600 GB
Replica set 9	3.4	32 vCPUs 128 GB	600 GB

ID	DB Version	DB Instance Class	Storage Space
Replica set 10	3.4	64 vCPUs 256 GB	600 GB
Replica set 11	4.0	2 vCPUs 4 GB	600 GB
Replica set 12	4.0	2 vCPUs 8 GB	600 GB
Replica set 13	4.0	4 vCPUs 8 GB	600 GB
Replica set 14	4.0	4 vCPUs 16 GB	600 GB
Replica set 15	4.0	8 vCPUs 16 GB	600 GB
Replica set 16	4.0	8 vCPUs 32 GB	600 GB
Replica set 17	4.0	16 vCPUs 32 GB	600 GB
Replica set 18	4.0	16 vCPUs 64 GB	600 GB
Replica set 19	4.0	32 vCPUs 128 GB	600 GB
Replica set 20	4.0	64 vCPUs 256 GB	600 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](#).
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', 'Spwd')
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 2-3 Service model

Service Model No.	Service Model
S1	100% insert
S2	90% update ,10% read

Service Model No.	Service Model
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
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

Specifications	Low-Level Data Volume	High-Level Data Volume
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.

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](#).
- Operations per Second (OPS): number of operations executed per second by a database
- In different service models, 10 million rows of data are preset for DB instances 3.4 with different shard specifications. For details about the test OPS, see the content in bold in [Table 2-5](#).

Table 2-5 Weak consistency: Test OPS of the database with 10 million rows of data.

shard Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
2 vCPUs 4 GB	20546	18399	14004	15564	17802	10707
2 vCPUs 8 GB	21070	30860	16188	19132	22058	10861
4 vCPUs 8 GB	38274	48373	16896	21819	23210	11176

4 vCPUs 16 GB	46871	79785	32461	74824	44682	21172
8 vCPUs 16 GB	67178	101073	33648	108414	50272	22424
8 vCPUs 32 GB	69476	111216	59425	110668	79308	36546
16 vCPUs 32 GB	72242	116774	62416	116201	83223	38373
16 vCPUs 64 GB	78146	115603	65235	123473	87620	46229
32 vCPUs 64 GB	81271	120227	67144	126411	91124	46078
32 vCPUs 128 GB	100947	150002	77951	147812	104516	55456
64 vCPUs 128 GB	108013	160502	83407	157158	112632	58737
64 vCPUs 256 GB	130382	192911	93939	192024	136045	60612

- In different service models, 100 million rows of data are preset for DB instances 3.4 with different shard specifications. For details about the test OPS, see the content in bold in [Table 2-6](#).

Table 2-6 Weak consistency: Test OPS of the database with 100 million rows of data.

shard Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
2 vCPUs 4 GB	8709	17810	13717	15208	11540	8534
2 vCPUs 8 GB	11384	22770	14892	20402	14809	9621
4 vCPUs 8 GB	18510	22912	15082	20488	14820	12098
4 vCPUs 16 GB	20982	23130	15612	20557	19209	12839

8 vCPUs 16 GB	28906	23642	16396	20614	19599	12917
8 vCPUs 32 GB	36881	26361	17059	21012	19647	13452
16 vCPUs 32 GB	38735	27679	17411	22062	20423	14124
16 vCPUs 64 GB	52220	39560	17493	29034	22123	13553
32 vCPUs 64 GB	54308	41142	17192	30095	22507	14395
32 vCPUs 128 GB	67345	50946	18348	37578	29329	17082
64 vCPUs 128 GB	71059	53512	19432	40318	31782	18277
64 vCPUs 256 GB	87877	66258	22917	47367	37631	22841

- In different service models, 10 million rows of data are preset for DB instances 4.0 with different shard specifications. For details about the test OPS, see the content in bold in [Table 2-7](#).

Table 2-7 Weak consistency: Test OPS of the database with 10 million rows of data.

shard Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
2 vCPUs 4 GB	19930	18031	13654	15268	17482	10471
2 vCPUs 8 GB	20438	30243	15783	18768	21661	10622
4 vCPUs 8 GB	37126	47406	16474	21404	22792	10930
4 vCPUs 16 GB	45465	78189	31649	73402	43878	20706
8 vCPUs 16 GB	65163	99052	32807	106354	49367	21931

8 vCPUs 32 GB	67392	108992	57939	108565	77880	35742
16 vCPUs 32 GB	70661	115441	60135	113993	81174	37629
16 vCPUs 64 GB	75802	113291	63604	121127	86043	45212
32 vCPUs 64 GB	78834	117822	66148	125972	89484	47020
32 vCPUs 128 GB	97919	147002	76002	145004	102635	54236
64 vCPUs 128 GB	105773	159292	82722	159154	113819	57532
64 vCPUs 256 GB	126471	189053	91591	188376	133596	59279

- In different service models, 100 million rows of data are preset for DB instances 4.0 with different shard specifications. For details about the test OPS, see the content in bold in [Table 2-8](#).

Table 2-8 Weak consistency: Test OPS of the database with 100 million rows of data.

shard Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
2 vCPUs 4 GB	8448	17454	13374	14919	11332	8346
2 vCPUs 8 GB	11042	22315	14520	20014	14542	9409
4 vCPUs 8 GB	17955	22454	14705	20099	14553	11832
4 vCPUs 16 GB	20353	22667	15222	20166	18863	12557
8 vCPUs 16 GB	28039	23169	15986	20222	19246	12633
8 vCPUs 32 GB	35775	25834	16633	20613	19293	13156

16 vCPUs 32 GB	37163	26725	17464	21643	20657	13813
16 vCPUs 64 GB	50653	38769	17056	28482	21725	13255
32 vCPUs 64 GB	52679	40319	17738	29621	22594	13785
32 vCPUs 128 GB	65325	49927	17889	36864	28801	16706
64 vCPUs 128 GB	71397	52421	18941	39414	31817	18675
64 vCPUs 256 GB	85241	64933	22344	46467	36954	22338

2.2.2 Replica Set

NOTE

- For details about the service model corresponding to the service model number, see [Table 2-3](#).
- Operations per Second (OPS): number of operations executed per second by a database
- In different service models, 10 million rows of data are preset for DB instances 3.4 with different specifications. For details about the test OPS, see the content in bold in [Table 2-9](#).

Table 2-9 Weak consistency: Test OPS of the database with 10 million rows of data.

DB Instance Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
2 vCPUs 4 GB	14826	9307	7412	9821	10886	5887
2 vCPUs 8 GB	17617	9337	8332	10274	11000	5676
4 vCPUs 8 GB	32826	10062	9728	10421	11144	5961
4 vCPUs 16 GB	38165	25248	9676	14830	14438	6846

8 vCPUs 16 GB	62880	52289	10319	15378	15282	6870
8 vCPUs 32 GB	63972	99628	21291	94992	54213	13028
16 vCPUs 32 GB	67570	104609	22155	98741	56523	13679
16 vCPUs 64 GB	71114	117312	56314	117787	86845	32131
32 vCPUs 64 GB	73758	121004	57561	121468	90316	33415
32 vCPUs 128 GB	92089	150876	73388	153121	112827	41769
64 vCPUs 128 GB	95672	156411	76123	158245	118140	43431
64 vCPUs 256 GB	118709	193819	95160	198786	146877	54300

- In different service models, 100 million rows of data are preset for DB instances 3.4 with different specifications. For details about the test OPS, see the content in bold in [Table 2-10](#).

Table 2-10 Weak consistency: Test OPS of the database with 100 million rows of data.

DB Instance Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
2 vCPUs 4 GB	7804	8569	6444	7231	7525	2550
2 vCPUs 8 GB	9931	8767	6744	8012	8457	4326
4 vCPUs 8 GB	15921	8799	7316	8835	8835	4392
4 vCPUs 16 GB	21050	12011	9141	10800	13352	4414
8 vCPUs 16 GB	28897	12514	9475	10855	13492	5184

8 vCPUs 32 GB	40044	12612	9508	11445	13569	7472
16 vCPUs 32 GB	42046	13242	9983	12017	14247	7845
16 vCPUs 64 GB	44178	12624	9601	11529	13612	7788
32 vCPUs 64 GB	45645	13125	9785	11570	14256	8091
32 vCPUs 128 GB	57408	16417	12481	14984	17695	10124
64 vCPUs 128 GB	61326	17613	13457	15732	18956	10799
64 vCPUs 256 GB	74598	21332	16203	19482	23005	13161

- In different service models, 10 million rows of data are preset for DB instances 4.0 with different specifications. For details about the test OPS, see the content in bold in [Table 2-11](#).

Table 2-11 Weak consistency: Test OPS of the database with 10 million rows of data.

DB Instance Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
2 vCPUs 4 GB	14381	9121	7227	9634	10690	5757
2 vCPUs 8 GB	17088	9150	8124	10079	10802	5551
4 vCPUs 8 GB	31841	9861	9485	10223	10943	5830
4 vCPUs 16 GB	37020	24743	9434	14548	14178	6695
8 vCPUs 16 GB	60994	51243	10061	15086	15007	6719
8 vCPUs 32 GB	62053	97635	20759	93187	53237	12741

16 vCPUs 32 GB	66155	102616	22776	97846	55871	13458
16 vCPUs 64 GB	68981	114966	54906	115549	85282	31424
32 vCPUs 64 GB	71740	119564	57102	130170	88693	32680
32 vCPUs 128 GB	89326	147858	71553	150212	110796	40850
64 vCPUs 128 GB	95578	158208	76561	162726	121851	43709
64 vCPUs 256 GB	115148	189943	92781	195009	144233	53105

- In different service models, 100 million rows of data are preset for DB instances 4.0 with different specifications. For details about the test OPS, see the content in bold in [Table 2-12](#).

Table 2-12 Weak consistency: Test OPS of the database with 100 million rows of data.

DB Instance Class	Service Model No.					
	s1	s6	s5	s4	s3	s2
2 vCPUs 4 GB	7570	8398	6283	7094	7390	2494
2 vCPUs 8 GB	9633	8592	6575	7860	8305	4231
4 vCPUs 8 GB	15443	8623	7133	8667	8676	4295
4 vCPUs 16 GB	20419	11771	8912	10595	13112	4317
8 vCPUs 16 GB	28030	12264	9238	10649	13249	5070
8 vCPUs 32 GB	38843	12360	9270	11228	13325	7308

16 vCPUs 32 GB	40376	12854	9670	11237	13858	7600
16 vCPUs 64 GB	42853	12372	9361	11310	13367	7617
32 vCPUs 64 GB	44567	12866	9735	11762	13901	7921
32 vCPUs 128 GB	55686	16089	12169	14699	17376	9901
64 vCPUs 128 GB	59584	17215	13420	16227	18532	10594
64 vCPUs 256 GB	72360	20905	15798	19112	22591	12871