

Document Database Service

Performance White Paper

Issue 01
Date 2022-08-30



Copyright © Huawei Technologies Co., Ltd. 2023. 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.

Security Declaration

Vulnerability

Huawei's regulations on product vulnerability management are subject to "Vul. Response Process". For details about the policy, see the following website:<https://www.huawei.com/en/psirt/vul-response-process>
For enterprise customers who need to obtain vulnerability information, visit:<https://securitybulletin.huawei.com/enterprise/en/security-advisory>

Contents

1 Test Method	1
2 Test Data	5
2.1 General-Purpose DB Instances.....	5
2.1.1 Cluster.....	5
2.1.2 Replica Set.....	10
2.2 Enhanced II Instances.....	14
2.2.1 Cluster.....	15
2.2.2 Replica Set.....	21

1 Test Method

This section describes the performance test of general-purpose and enhanced II DB instances of DDS Community Edition 3.4, 4.0, 4.2, and 4.4, including the test environment and procedure.

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

Test Environment

- Region: AP-Singapore
- AZ: AZ1
- Elastic Cloud Server (ECS):
 - For a general-purpose DB instance, select an ECS with s3.2xlarge.2 flavor with 8 vCPUs, 16 GB of memory, and CentOS7.5 64 bit image.
 - For an enhanced II DB instance, select an ECS with c3.2xlarge.2 flavor with 8 vCPUs, 16 GB of memory, and CentOS 7.6 x86_64 bit image.
- DB instances to be tested:
 - The cluster instance has two mongos nodes and two shard nodes, and the storage space is 600 GB.
 - The replica set instance has three nodes, and the storage space is 600 GB.

Test Tool

YCSB is an open-source database performance test tool. In this test, YCSB [0.17.0](#) 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. Use YCSB to run the following command to load the test data:

```
./bin/ycsb load mongodb -s -p workload=site.ycsb.workloads.CoreWorkload -p recordcount=$
```

```
{recordcount} -p mongodb.url="mongodb://{userName}:{password}@{mongosIP}:{port}/ycsb?authSource=admin" -threads ${threadNum}
```

 NOTE

- **recordcount=\${recordcount}** indicates the total amount of data loaded to the DDS instance.
 - **mongodb.url="mongodb://{userName}:{password}@{mongosIP}:{port}/ycsb?authSource=admin"** indicates the HA connection address of the DDS instance.
 - **threads \${threadNum}** indicates the number of concurrent threads on the client.
2. Use YCSB to run the following command to test the performance:

```
./bin/ycsb run mongodb -s -p workload=site.ycsb.workloads.CoreWorkload -p recordcount=${recordcount} -p operationcount=${operationcount} -p insertproportion=0 -p readproportion=1 -p updateproportion=0 -p requestdistribution=zipfian -p mongodb.url="mongodb://{userName}:{password}@{mongosIP}:{port}/ycsb?authSource=admin" -threads ${threadNum}
```

 NOTE

- **recordcount=\${recordcount}** indicates the total amount of data loaded to the DDS instance.
- **operationcount=\${operationcount}** indicates the total number of read/write operations.
- **insertproportion=0** indicates the proportion of data loading operations.
- **readproportion=1** indicates the proportion of read operations.
- **updateproportion=0** indicates the proportion of update operations.
- **mongodb.url="mongodb://{userName}:{password}@{mongosIP}:{port}/ycsb?authSource=admin"** indicates the HA connection address of the DDS instance.
- **threads \${threadNum}** indicates the number of concurrent threads on the client.
- The values of **recordcount** and **operationcount** are the same.
- The sum of the values of **insertproportion**, **readproportion**, and **updateproportion** is 1.

Testing Models

- Service model

Table 1-1 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

Service Model No.	Service Model
S7	95% read ,5% update
S8	5% read ,95% update
S9	100% update

- 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 1-2 Data volume to be prepared

Specifications	Low-Level Data Volume	High-Level Data Volume	Number of Concurrent Threads
1 vCPUs 4 GB	Records: 1,000,000	Records: 10,000,000	100
2 vCPUs 4 GB	Records: 1,000,000	Records: 10,000,000	100
2 vCPUs 8 GB	Records: 2,000,000	Records: 20,000,000	100
2 vCPUs 16 GB	Records: 4,000,000	Records: 40,000,000	100
4 vCPUs 8 GB	Records: 2,000,000	Records: 20,000,000	100
4 vCPUs 16 GB	Records: 4,000,000	Records: 40,000,000	100
4 vCPUs 32 GB	Records: 10,000,000	Records: 100,000,000	100
8 vCPUs 16 GB	Records: 4,000,000	Records: 40,000,000	100
8 vCPUs 32 GB	Records: 10,000,000	Records: 100,000,000	100
8 vCPUs 64 GB	Records: 10,000,000	Records: 100,000,000	100
16 vCPUs 32 GB	Records: 10,000,000	Records: 100,000,000	100

Specifications	Low-Level Data Volume	High-Level Data Volume	Number of Concurrent Threads
16 vCPUs 64 GB	Records: 10,000,000	Records: 100,000,000	100
16 vCPUs 128 GB	Records: 10,000,000	Records: 100,000,000	100
32 vCPUs 64 GB	Records: 10,000,000	Records: 100,000,000	200
32 vCPUs 128 GB	Records: 10,000,000	Records: 100,000,000	200
64 vCPUs 256 GB	Records: 10,000,000	Records: 100,000,000	400

- 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 Test Data

2.1 General-Purpose DB Instances

2.1.1 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, test the OPS of DB instances 3.4 with different shard classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-1](#).

Table 2-1 Weak consistency: OPS of DB instances with low-level data volume

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, test the OPS of DB instances 3.4 with different shard classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-2](#).

Table 2-2 Weak consistency: OPS of DB instances with high-level data volume

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, test the OPS of DB instances 4.0 with different shard classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-3](#).

Table 2-3 Weak consistency: OPS of DB instances with low-level data volume

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, test the OPS of DB instances 4.0 with different shard classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-4](#).

Table 2-4 Weak consistency: OPS of DB instances with high-level data volume

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

- In different service models, test the OPS of DB instances 4.2 with different shard classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-5](#).

Table 2-5 Weak consistency: OPS of DB instances with low-level data volume

shard Class	Service Model No.					
	s1	s6	s7	s5	s8	s9

1 vCPU 4 GB	9612	14234	13031	8308	5750	5495
2 vCPUs 4 GB	14751	20652	20027	14408	9164	8683
2 vCPUs 8 GB	15988	22472	21839	17474	13914	13886
4 vCPUs 8 GB	27347	45729	43174	29658	22421	22091
4 vCPUs 16 GB	28662	46096	43582	30020	22868	22936
8 vCPUs 16 GB	45312	81612	78355	55312	39526	38223
8 vCPUs 32 GB	45839	82340	78779	55991	39858	38714

- In different service models, test the OPS of DB instances 4.2 with different shard classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-6](#).

Table 2-6 Weak consistency: OPS of DB instances with high-level data volume

shard Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
1 vCPU 4 GB	7584	12741	11058	6355	4339	5363
2 vCPUs 4 GB	12929	20043	19063	10770	6633	6104
2 vCPUs 8 GB	14362	22294	20614	11827	7937	7783
4 vCPUs 8 GB	23889	39662	36583	25342	19729	18691
4 vCPUs 16 GB	25587	44711	41377	28728	22914	22139
8 vCPUs 16 GB	42111	73541	67942	47490	36676	35923
8 vCPUs 32 GB	42476	74210	68257	48045	37571	36665

- In different service models, test the OPS of DB instances 4.4 with different shard classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-7](#).

Table 2-7 Weak consistency: OPS of DB instances with low-level data volume

shard Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
1 vCPU 4 GB	10274	15235	13867	8323	5912	5784
2 vCPUs 4 GB	15244	21385	20561	14778	9436	8673
2 vCPUs 8 GB	16481	23205	22374	17844	14186	13876
4 vCPUs 8 GB	27457	45456	43357	30063	22456	22264
4 vCPUs 16 GB	27982	46049	43976	31867	22736	22476
8 vCPUs 16 GB	45541	81983	78721	55234	40019	38742
8 vCPUs 32 GB	46068	82712	79145	55913	40351	39233

- In different service models, test the OPS of DB instances 4.4 with different shard classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-8](#).

Table 2-8 Weak consistency: OPS of DB instances with high-level data volume

shard Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
1 vCPU 4 GB	8245	13741	11893	6369	4501	5651
2 vCPUs 4 GB	13421	20776	19597	11140	6904	6094
2 vCPUs 8 GB	14855	23026	21148	12196	8209	7773
4 vCPUs 8 GB	23999	39388	36766	25746	19763	18864
4 vCPUs 16 GB	24906	41098	38765	29843	21846	20674
8 vCPUs 16 GB	42340	73912	68308	47411	37169	36442
8 vCPUs 32 GB	42705	74581	68623	47967	38064	37184

2.1.2 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, test the OPS of DB instances 3.4 with different classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-9](#).

Table 2-9 Weak consistency: OPS of DB instances with low-level data volume

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, test the OPS of DB instances 3.4 with different classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-10](#).

Table 2-10 Weak consistency: OPS of DB instances with high-level data volume

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, test the OPS of DB instances 4.0 with different classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-11](#).

Table 2-11 Weak consistency: OPS of DB instances with low-level data volume

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, test the OPS of DB instances 4.0 with different classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-12](#).

Table 2-12 Weak consistency: OPS of DB instances with high-level data volume

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

- In different service models, test the OPS of DB instances 4.2 with different classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-13](#).

Table 2-13 Weak consistency: OPS of DB instances with low-level data volume

DB Instance Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
1 vCPU 4 GB	7057	12559	10200	5587	3274	3245
2 vCPUs 4 GB	14982	25847	21504	11247	6332	6151
2 vCPUs 8 GB	15013	26487	21553	11430	6900	6811
4 vCPUs 8 GB	26832	48422	39484	21875	13873	13990
4 vCPUs 16 GB	27119	48455	39882	22256	13968	14432

8 vCPUs 16 GB	46318	95619	78473	41528	28963	28522
8 vCPUs 32 GB	46631	97275	82517	43623	30883	29997

- In different service models, test the OPS of DB instances 4.2 with different classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-14](#).

Table 2-14 Weak consistency: OPS of DB instances with high-level data volume

DB Instance Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
1 vCPU 4 GB	4807	8665	7337	3646	2492	2915
2 vCPUs 4 GB	7254	17962	15065	6533	3465	3004
2 vCPUs 8 GB	7587	18261	15253	7756	5131	4894
4 vCPUs 8 GB	15577	32918	28480	14420	9575	8830
4 vCPUs 16 GB	22200	33115	29420	16973	11821	11308
8 vCPUs 16 GB	39838	58903	51712	32424	23810	23172
8 vCPUs 32 GB	40901	62003	54651	35205	25976	25053

- In different service models, test the OPS of DB instances 4.4 with different classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-15](#).

Table 2-15 Weak consistency: OPS of DB instances with low-level data volume

DB Instance Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
1 vCPU 4 GB	7374	12458	10065	5066	3569	3310

2 vCPUs 4 GB	15203	25941	21320	10984	6256	6023
2 vCPUs 8 GB	15001	25847	21985	11364	6546	6298
4 vCPUs 8 GB	27001	48569	39748	21238	14107	14023
4 vCPUs 16 GB	27647	48846	40215	22054	14235	14046
8 vCPUs 16 GB	46684	95451	78314	40244	29648	28672
8 vCPUs 32 GB	46997	97107	82358	42339	31569	30147

- In different service models, test the OPS of DB instances 4.4 with different classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-16](#).

Table 2-16 Weak consistency: OPS of DB instances with high-level data volume

DB Instance Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
1 vCPU 4 GB	5064	8813	7436	3768	3051	3028
2 vCPUs 4 GB	7655	18344	15542	6456	3634	3312
2 vCPUs 8 GB	7988	18643	15730	7680	5299	5202
4 vCPUs 8 GB	15876	33645	28546	14239	9234	8768
4 vCPUs 16 GB	18349	33842	29486	16793	11480	11246
8 vCPUs 16 GB	40647	60238	54536	32344	25673	22351
8 vCPUs 32 GB	41710	62346	56432	35124	27839	24231

2.2 Enhanced II Instances

2.2.1 Cluster

 NOTE

- For details about the service model corresponding to the service model number, see [Service model](#).
- Operations per Second (OPS): number of operations executed per second by a database
- In different service models, test the OPS of DB instances 3.4 with different shard classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-17](#).

Table 2-17 Weak consistency: OPS of DB instances with low-level data volume

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, test the OPS of DB instances 3.4 with different shard classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-18](#).

Table 2-18 Weak consistency: OPS of DB instances with high-level data volume

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, test the OPS of DB instances 4.0 with different shard classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-19](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-19 Weak consistency: OPS of DB instances with low-level data volume

shard Class	Service Model No.			
	s1	s6	s7	s5
2 vCPUs 4 GB	18236	28013	26529	18907
2 vCPUs 8 GB	16942	27931	25365	18706
4 vCPUs 8 GB	25979	39312	37286	26043
4 vCPUs 16 GB	36576	53862	51511	36715
8 vCPUs 16 GB	49335	73368	68763	48786
8 vCPUs 32 GB	66986	99692	92745	67052
16 vCPUs 32 GB	82218	109324	106841	85882
16 vCPUs 64 GB	100293	134211	131163	105551
32 vCPUs 64 GB	144640	185165	181450	154899
32 vCPUs 128 GB	141616	186344	180511	156879
64 vCPUs 256 GB	163832	221621	242273	245718

- In different service models, test the OPS of DB instances 4.0 with different shard classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-20](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-20 Weak consistency: OPS of DB instances with high-level data volume

shard Class	Service Model No.			
	s1	s6	s7	s5
2 vCPUs 4 GB	14622	25400	21966	14243
2 vCPUs 8 GB	15029	25896	22654	14376
4 vCPUs 8 GB	21695	31590	27398	18865
4 vCPUs 16 GB	31906	50302	44697	27710
8 vCPUs 16 GB	43145	62259	55496	33406
8 vCPUs 32 GB	56384	77832	70588	35539
16 vCPUs 32 GB	69135	92815	83404	46797
16 vCPUs 64 GB	90531	123229	114717	57297
32 vCPUs 64 GB	139652	170371	161543	73667
32 vCPUs 128 GB	142140	181753	177054	111161
64 vCPUs 256 GB	177758	252032	249090	252073

- In different service models, test the OPS of DB instances 4.2 with different shard classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-21](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-21 Weak consistency: OPS of DB instances with low-level data volume

shard Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
2 vCPUs 8 GB	16400	25734	24121	17383	13721	13475
2 vCPUs 16 GB	16489	25789	24283	17417	13765	13608
4 vCPUs 16 GB	31904	49090	47196	35078	26331	26519
4 vCPUs 32 GB	32150	49328	47227	35781	26395	26720

8 vCPUs 32 GB	53819	90702	85896	62155	46241	45040
8 vCPUs 64 GB	54753	93158	86468	62868	46343	45804
16 vCPUs 64 GB	74752	132352	124941	92073	68398	65499
16 vCPUs 128 GB	74843	134285	129159	93630	68866	66104

- In different service models, test the OPS of DB instances 4.2 with different shard classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-22](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-22 Weak consistency: OPS of DB instances with high-level data volume

shard Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
2 vCPUs 8 GB	15219	23720	21281	12014	8337	7843
2 vCPUs 16 GB	15455	24551	21676	12943	9312	8575
4 vCPUs 16 GB	27189	42101	39332	25368	20222	20026
4 vCPUs 32 GB	28334	45196	42295	30231	23794	23091
8 vCPUs 32 GB	50561	81715	75940	54235	42504	41380
8 vCPUs 64 GB	51195	86148	77881	54814	42568	42350
16 vCPUs 64 GB	70625	117463	110415	79770	61999	59432
16 vCPUs 128 GB	72678	128983	126588	88769	65706	61499

- In different service models, test the OPS of DB instances 4.4 with different shard classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-23](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-23 Weak consistency: OPS of DB instances with low-level data volume

shard Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
2 vCPUs 8 GB	16869	25132	23956	18025	14025	13986
2 vCPUs 16 GB	16859	26025	24986	18025	14025	14005
4 vCPUs 16 GB	32005	48952	47965	35968	27002	27025
4 vCPUs 32 GB	32968	50235	47962	35268	26541	26531
8 vCPUs 32 GB	54003	91025	86023	62538	46895	45102
8 vCPUs 64 GB	55012	93458	87126	62325	46982	46005
16 vCPUs 64 GB	75032	135962	125365	92986	68321	65823
16 vCPUs 128 GB	75125	135102	130802	94206	69035	66423

- In different service models, test the OPS of DB instances 4.4 with different shard classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-24](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-24 Weak consistency: OPS of DB instances with high-level data volume

shard Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
2 vCPUs 8 GB	15396	23568	21689	11986	8386	7953
2 vCPUs 16 GB	16025	25125	21932	13058	9365	8602

4 vCPUs 16 GB	27869	42698	39125	26015	20963	19996
4 vCPUs 32 GB	28964	45962	43069	30954	34025	23855
8 vCPUs 32 GB	50786	82456	76025	53214	42586	41865
8 vCPUs 64 GB	51962	86232	77962	54862	42968	42059
16 vCPUs 64 GB	70862	112582	118685	80056	62008	60487
16 vCPUs 128 GB	71963	128642	120569	87256	65754	61985

2.2.2 Replica Set

NOTE

- For details about the service model corresponding to the service model number, see [Service model](#).
- Operations per Second (OPS): number of operations executed per second by a database
- In different service models, test the OPS of DB instances 3.4 with different classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-25](#).

Table 2-25 Weak consistency: OPS of DB instances with low-level data volume

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, test the OPS of DB instances 3.4 with different classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-26](#).

Table 2-26 Weak consistency: OPS of DB instances with high-level data volume

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, test the OPS of DB instances 4.0 with different classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-27](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-27 Weak consistency: OPS of DB instances with low-level data volume

DB Instance Class	Service Model No.			
	s1	s6	s7	s5
2 vCPUs 4 GB	11902	27719	21480	11056
2 vCPUs 8 GB	11807	29933	22895	11449
4 vCPUs 8 GB	17584	40504	31579	16101
4 vCPUs 16 GB	24560	60560	46787	23382
8 vCPUs 16 GB	34262	77401	64056	32177
8 vCPUs 32 GB	43510	112973	91838	40838
16 vCPUs 32 GB	68944	150944	126309	58089

16 vCPUs 64 GB	93615	221408	181915	85458
32 vCPUs 64 GB	105738	261882	206798	109913
32 vCPUs 128 GB	115401	260751	231041	104299
64 vCPUs 256 GB	118546	295335	274551	281141

- In different service models, test the OPS of DB instances 4.0 with different classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-28](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-28 Weak consistency: OPS of DB instances with high-level data volume

DB Instance Class	Service Model No.			
	s1	s6	s7	s5
2 vCPUs 4 GB	9566	15988	13458	8209
2 vCPUs 8 GB	7476	17263	13988	8559
4 vCPUs 8 GB	15202	19333	16110	9996
4 vCPUs 16 GB	22163	23924	22842	13052
8 vCPUs 16 GB	30065	31584	29172	16850
8 vCPUs 32 GB	37637	31627	29464	17004
16 vCPUs 32 GB	51973	39088	36279	21907
16 vCPUs 64 GB	66750	50288	46154	25441
32 vCPUs 64 GB	87285	66944	61949	34386
32 vCPUs 128 GB	94519	102497	91120	38757
64 vCPUs 256 GB	119649	312685	301940	301425

- In different service models, test the OPS of DB instances 4.2 with different classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-29](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-29 Weak consistency: OPS of DB instances with low-level data volume

DB Instance Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
2 vCPUs 8 GB	14812	25271	20558	10798	7118	6403
2 vCPUs 16 GB	14990	25303	21272	11184	7369	6508
4 vCPUs 16 GB	26430	46584	39433	22286	14769	14650
4 vCPUs 32 GB	27161	49916	41301	23402	16545	15921
8 vCPUs 32 GB	47675	90022	78644	43393	30949	30053
8 vCPUs 64 GB	48008	94013	80441	44415	31473	30422
16 vCPUs 64 GB	72842	173838	143535	76799	51723	48473
16 vCPUs 128 GB	78344	180946	148867	79640	53445	51666

- In different service models, test the OPS of DB instances 4.2 with different classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-30](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-30 Weak consistency: OPS of DB instances with high-level data volume

DB Instance Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
2 vCPUs 8 GB	9539	16597	13886	7547	4913	4607

2 vCPUs 16 GB	10763	16713	14714	8019	5768	5290
4 vCPUs 16 GB	23682	32746	28598	17193	11168	10072
4 vCPUs 32 GB	24709	37024	32043	19545	12881	11718
8 vCPUs 32 GB	43609	64288	58465	37392	27209	26417
8 vCPUs 64 GB	43735	67299	59211	37989	27396	26696
16 vCPUs 64 GB	67483	115287	100895	61909	42337	40532
16 vCPUs 128 GB	74203	142301	122038	69775	46934	45159

- In different service models, test the OPS of DB instances 4.4 with different classes based on the preset low-level data volume in 2. For details, see the content in bold in [Table 2-31](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-31 Weak consistency: OPS of DB instances with low-level data volume

DB Instance Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
2 vCPUs 8 GB	15426	26148	20498	11256	7065	6321
2 vCPUs 16 GB	15002	26004	20869	11598	7405	6591
4 vCPUs 16 GB	25984	48521	40154	21582	15025	15602
4 vCPUs 32 GB	27589	50036	40986	24015	15985	16053
8 vCPUs 32 GB	46855	89952	77256	43856	30120	29935
8 vCPUs 64 GB	47925	95426	80986	45012	31859	30998
16 vCPUs 64 GB	73568	178690	139562	75986	52006	49856

16 vCPUs 128 GB	78652	179526	151509	80024	53798	52486
----------------------	--------------	---------------	---------------	--------------	--------------	--------------

- In different service models, test the OPS of DB instances 4.4 with different classes based on the preset high-level data volume in 2. For details, see the content in bold in [Table 2-32](#).

 **NOTE**

The table lists only the performance data in typical scenarios.

Table 2-32 Weak consistency: OPS of DB instances with high-level data volume

DB Instance Class	Service Model No.					
	s1	s6	s7	s5	s8	s9
2 vCPUs 8 GB	9602	16985	13954	7601	5012	4682
2 vCPUs 16 GB	11542	17695	14856	8108	5768	5320
4 vCPUs 16 GB	24568	31986	28453	17869	11586	11520
4 vCPUs 32 GB	25014	37985	32986	20458	13054	11586
8 vCPUs 32 GB	43598	68523	59025	37892	27854	27654
8 vCPUs 64 GB	43869	68514	60254	38053	27856	26985
16 vCPUs 64 GB	68526	118960	108965	62045	42398	41582
16 vCPUs 128 GB	75684	145268	128964	70058	47025	46054