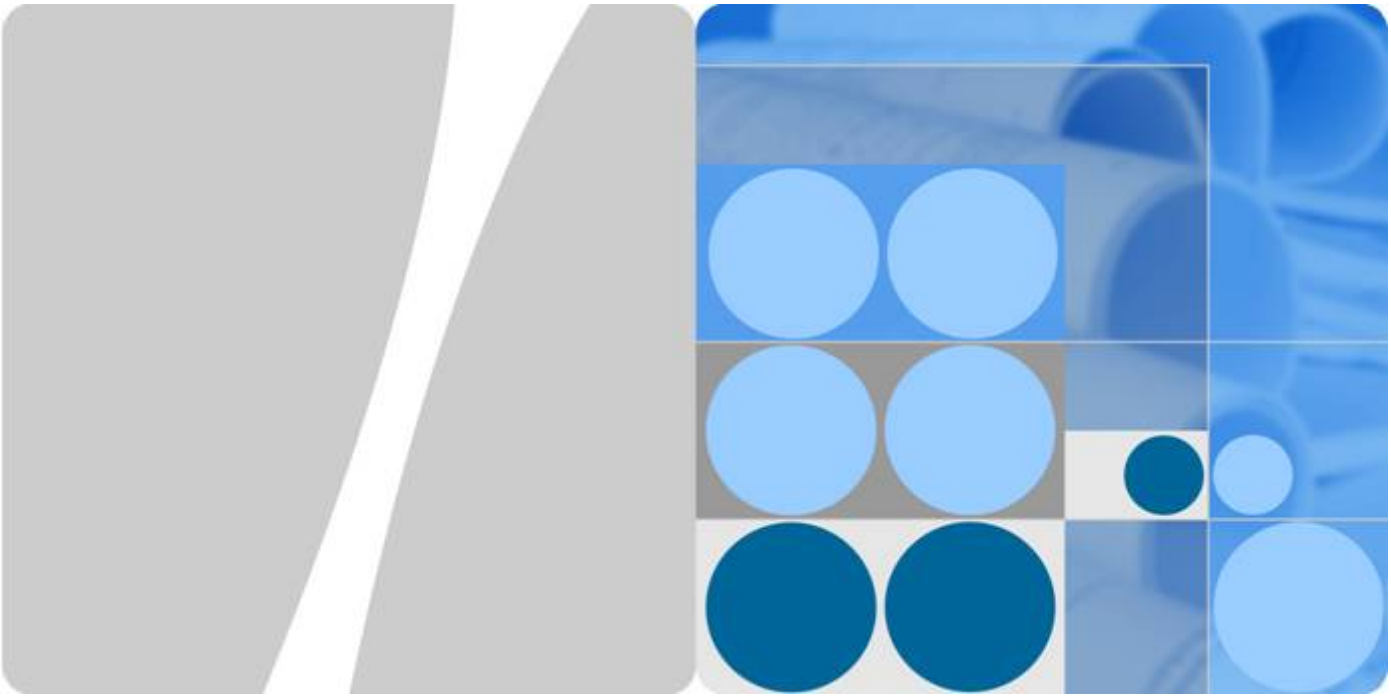

GaussDB NoSQL

Performance White Paper

Issue 01

Date 2020-12-10

HUAWEI TECHNOLOGIES CO., LTD.

HUAWEI

Copyright © Huawei Technologies Co., Ltd. 2020. 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


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 GaussDB(for Cassandra)..... 1
1.1 Test Method..... 1
1.2 Test Data..... 3
2 GaussDB(for Mongo)..... 4
2.1 Cluster..... 4
2.1.1 Test Method..... 4
2.1.2 Test Data..... 7
2.2 Replica Set..... 9
2.2.1 Test Method..... 9
2.2.2 Test Data..... 12

 Copyright © Huawei Technologies Co., Ltd.

1 GaussDB(for Cassandra)

1.1 Test Method

This section describes the performance test of GaussDB(for Cassandra) instances, including the test environment, procedures, and results.

Test Environment

- Region: AP-Singapore
- AZ: AZ1, AZ2, AZ3 (three-AZ deployment)
- Elastic Cloud Server (ECS): h3.4xlarge.2 with 16 vCPUs, 32 GB of memory, and CentOS 7.5 64 bit image.
- Each DB instance includes three nodes.
- The test covers the following instance classes. For details, see **Table 1-1**.

Table 1-1 DB instance classes

No.	Instance Class
cluster1	4 vCPUs | 16 GB
cluster2	8 vCPUs | 32 GB
cluster3	16 vCPUs | 64 GB
cluster4	32 vCPUs | 128 GB

Test Tool

YCSB is an open-source database performance test tool. In this test, YCSB 0.15.0 is used.

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

 Copyright © Huawei Technologies Co., Ltd.

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**, **updateproportion**, **scanproportion**, and **readmodifywriteproportion** in the **workload** file by referring to **Table 1-2**.

   Set the value of **recordcount** in the **workload** file by referring to **Table 1-3**.

2. Use workload-insert-mostly as an example. Run the following command to prepare test data:

   sh bin/ycsb.sh load cassandra-cql -P workloads/workload-insert-mostly -p "hosts=${ContactPoints}" -p "port=${port}" -p "cassandra.username=${username}" -p "cassandra.password=${password}" -p operationcount=400000000 -p recordcount=400000000 -p exportfile=./data_load.exp -threads ${threadNum} -s > data_load.log 2>&1 &

3. Use workload-insert-mostly as an example. Run the following command to test the performance:

   sh bin/ycsb.sh run cassandra-cql -P workloads/workload-insert-mostly -p "hosts=${ContactPoints}" -p "port=${port}" -p "cassandra.username=${username}" -p "cassandra.password=${password}" -p operationcount=9000000 -p recordcount=9000000 -p maxexecutiontime=3600 -p exportfile=./workload-insert-mostly.exp -threads ${threadNum} -s > workload-insert-mostly_run.log 2>&1 &

Testing Models

- Workload model

  Table 1-2 Test model

No.	Testing Model
workload-read-mostly	95% read, 5% update
workload-read-write-combination	50% update, 50% read
workload-read-modify-write	50% read, 50% readmodifywrite
workload-mixed-operational-analytical	65% read, 25% update, 10% insert
workload-insert-mostly	90% insert, 10% read

- Data model

  fieldlength=100, fieldcount=10

- Data volume to be prepared

  Test the DB instances of different classes using the following data volume.

 Copyright © Huawei Technologies Co., Ltd.

The following table lists the data volume.

Table 1-3 Data volume to be prepared

No.	Instance Class	Data volume to be prepared
cluster1	4 vCPUs | 16 GB	50 GB
cluster2	8 vCPUs | 32 GB	100 GB
cluster3	16 vCPUs | 64 GB	200 GB
cluster4	32 vCPUs | 128 GB	400 GB

- Test mode

  In the test, high-performance mode is used.

1.2 Test Data

NOTE

For details about the test model corresponding to the test model number, see **Table 1-2**.

Test OPS in different service models based on DB instance class. For details, see the content in bold in **Table 1-4**.

Table 1-4 Test data

High performance mode: Test OPS based on the prepared data volume					
Node Class		4 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32 vCPUs | 128 GB
Concurrent client requests		32	64	128	256
Data volume to be prepared		50 GB	100 GB	200 GB	400 GB
Test Model Number	workload-read-mostly	**15627**	**44612**	**90713**	**169037**
	workload-read-write-combination	**19565**	**46240**	**95065**	**158793**
	workload-read-modify-write	**11768**	**29488**	**59332**	**96964**
	workload-mixed-operational-analytical	**17534**	**40214**	**80661**	**137616**
	workload-insert-mostly	**23830**	**44548**	**67290**	**111540**

 Copyright © Huawei Technologies Co., Ltd.

2 GaussDB(for Mongo)

2.1 Cluster

2.1.1 Test Method

This section describes the performance test of GaussDB(for Mongo) cluster 3.4, including the test environment, procedures, and results.

Test Environment

- Region: CN North-Beijing4
- 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 DB instance includes two shards.
- For details about all the DB instance classes, see **Table 2-1**.

Table 2-1 Cluster instance class

Cluster ID	shard Class
Cluster 1	4 vCPUs | 16 GB
Cluster 2	8 vCPUs | 32 GB
Cluster 3	16 vCPUs | 64 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**.

 Copyright © Huawei Technologies Co., Ltd.

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 2-2**.

   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

Testing Models

- Workload model

  Table 2-2 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
S7	95% read, 5% update
S8	10% read, 90% insert
S9	5% read, 95% update

- Number of concurrent threads

 Copyright © Huawei Technologies Co., Ltd.

Table 2-3 Number of concurrent threads

shard Class	Concurrent Threads of the Client
4 vCPUs | 16 GB	32
8 vCPUs | 32 GB	64
16 vCPUs | 64 GB	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

Cluster ID	shard Class	Low-Level Data Volume	High-Level Data Volume
Cluster 1	4 vCPUs | 16 GB	Storage space: 200 GB Record counts: 200,000,000	Storage space: 1000 GB Record counts: 1,000,000,000
Cluster 2	8 vCPUs | 32 GB	Storage space: 400 GB Record counts: 400,000,000	Storage space: 2000 GB Record counts: 2,000,000,000
Cluster 3	16 vCPUs | 64 GB	Storage space: 800 GB Record counts: 800,000,000	Storage space: 4000 GB Record counts: 4,000,000,000

- Data consistency model

  This test involves the following data consistency models:

  – 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.

  – Strong consistency: For the write concern setting of **{w: 3, j: true}**, no acknowledgment is returned until data is written to disks on the three nodes.

 Copyright © Huawei Technologies Co., Ltd.

2.1.2 Test Data

Weak Consistency

- Test OPS in different test service models based on different shard classes and the preset low-level data volume. For details, see the content in bold in **Table 2-5**.

Table 2-5 Test data

Weak consistency: Test the OPS of the database with low-level data volume.				
shard Class		4 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB
Data volume to be prepared		200 GB	400 GB	800 GB
Test service model	100% I	**22837**	**38823**	**66000**
	90% U 10% R	**11944**	**20305**	**34518**
	65% R 10% U 25% I	**14291**	**24295**	**41301**
	90% R 5% I 5% U	**13754**	**23382**	**39749**
	50% R 50% U	**12622**	**21457**	**36477**
	100% R	**15283**	**25981**	**44168**
	95% R 5% U	**13841**	**23529**	**40000**
	10% R 90% I	**16343**	**27783**	**47231**
	5% R 95% U	**11973**	**20354**	**34601**

- Test OPS in different test service models based on different shard classes and the preset high-level data volume. For details, see the content in bold in **Table 2-6**.

Table 2-6 Test data

Weak consistency: Test the OPS of the database with high-level data volume.				
shard Class		4 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB
Data volume to be prepared		1000 GB	2000 GB	4000 GB
Test service model	100% I	**19608**	**33334**	**56668**

 Copyright © Huawei Technologies Co., Ltd.

Weak consistency: Test the OPS of the database with high-level data volume.				
	90% U 10% R	**10998**	**18697**	**31786**
	65% R 10% U 25% I	**13260**	**22542**	**38322**
	90% R 5% I 5% U	**13185**	**22414**	**38104**
	50% R 50% U	**11869**	**20177**	**34301**
	100% R	**15531**	**26403**	**44885**
	95% R 5% U	**12898**	**21926**	**37275**
	10% R 90% I	**19708**	**33504**	**56956**
	5% R 95% U	**11845**	**20137**	**34233**

Strong Consistency

- Test OPS in different test service models based on different shard classes and the preset low-level data volume. For details, see the content in bold in **Table 2-7**.

Table 2-7 Test data

Strong consistency: Test the OPS of the database with low-level data volume.				
shard Class		4 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB
Data volume to be prepared		200 GB	400 GB	800 GB
Test service model	100% I	**9571**	**16271**	**27660**
	90% U 10% R	**7402**	**12583**	**21391**
	65% R 10% U 25% I	**10968**	**18645**	**31696**
	90% R 5% I 5% U	**12767**	**21704**	**36897**
	50% R 50% U	**9473**	**16105**	**27378**
	100% R	**15287**	**25988**	**44180**
	95% R 5% U	**13338**	**22674**	**38546**
	10% R 90% I	**9018**	**15331**	**26063**

 Copyright © Huawei Technologies Co., Ltd.

Strong consistency: Test the OPS of the database with low-level data volume.				
	5% R 95% U	**7266**	**12351**	**20997**

- Test OPS in different test service models based on different shard classes and the preset high-level data volume. For details, see the content in bold in **Table 2-8**.

Table 2-8 Test data

Strong consistency: Test the OPS of the database with high-level data volume.				
shard Class		4 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB
Data volume to be prepared		1000 GB	2000 GB	4000 GB
Test service model	100% I	**7293**	**12398**	**21076**
	90% U 10% R	**6608**	**11234**	**19099**
	65% R 10% U 25% I	**9844**	**16735**	**28449**
	90% R 5% I 5% U	**10993**	**18688**	**31769**
	50% R 50% U	**8391**	**14265**	**24251**
	100% R	**13208**	**22454**	**38172**
	95% R 5% U	**11206**	**19051**	**32386**
	10% R 90% I	**9166**	**15583**	**26490**
	5% R 95% U	**6330**	**10761**	**18294**

2.2 Replica Set

2.2.1 Test Method

This section describes the performance test of GaussDB(for Mongo) replica set 4.0, including the test environment, procedures, and results.

Test Environment

- Region: CN North-Beijing4
- AZ: AZ1

 Copyright © Huawei Technologies Co., Ltd.

- Elastic Cloud Server (ECS): s3.2xlarge.2 flavor with 8 vCPUs, 16 GB of memory, and CentOS 7.5 64 bit image.
- Replica set instance specifications: All instance specifications are tested. For details, see **Table 2-9**.

Table 2-9 Replica set specifications

Replica Set ID	DB Version	Instance Class
Replica set 1	4.0	4 vCPUs | 16 GB
Replica set 2	4.0	8 vCPUs | 32 GB
Replica set 3	4.0	16 vCPUs | 64 GB
Replica set 4	4.0	32 vCPUs | 128 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 2-10**.

   Set the value of **recordcount** in the **workload** file by referring to **Table 2-11**.

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}@${mongodIP}:${port},${mongodIP}:${port},${mongodIP}:${port}/ycsb?authSource=admin&replicaSet=replica&readPreference=secondary -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}@${mongodIP}:${port},${mongodIP}:${port},${mongodIP}:${port}/ycsb?authSource=admin&replicaSet=replica&readPreference=secondary -threads ${threadNum} -p maxexecutiontime=1800 1>workload_s1_run.result 2> workload_s1_run.log

 Copyright © Huawei Technologies Co., Ltd.

NOTE

- ${mongodIP} indicates the private IP addresses of the primary and secondary nodes in the GaussDB(for Mongo) replica set instance.
- ${password} is the administrator password of the GaussDB(for Mongo) replica set 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 2-10 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 replica set instance.

  For details, see the following table.

  Table 2-11 Data volume to be prepared

Replica Set ID	Instance Class	Low-Level Data Volume	High-Level Data Volume
Replica set 1	4 vCPUs | 16 GB	Storage space: 10 GB RecordCount: 10000000	Storage space: 100 GB Record counts: 100,000,000
Replica set 2	8 vCPUs | 32 GB	Storage space: 10 GB RecordCount: 10000000	Storage space: 100 GB Record counts: 100,000,000

 Copyright © Huawei Technologies Co., Ltd.

Replica Set ID	Instance Class	Low-Level Data Volume	High-Level Data Volume
Replica set 3	16 vCPUs | 64 GB	Storage space: 10 GB RecordCount: 10000000	Storage space: 100 GB Record counts: 100,000,000
Replica set 4	32 vCPUs | 128 GB	Storage space: 10 GB RecordCount: 10000000	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.

2.2.2 Test Data

NOTE

For details about the service model corresponding to the service model number, see **Table 2-10**.

- Test OPS in different test service models based on different instance classes and the preset 10 million rows of data. For details, see the content in bold in **Table 2-12**.

Table 2-12 Test data

Weak consistency: Test OPS of the database with 10 million rows of data.					
Instance Class		4 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32 vCPUs | 128 GB
No.	s6	**29286**	**64882**	**97377**	**12107**
	s5	**16829**	**35224**	**60616**	**81019**
	s4	**23338**	**75905**	**89156**	**101306**
	s3	**19959**	**55442**	**79135**	**98611**
	s2	**13399**	**26732**	**39276**	**50833**
	s1	**35263**	**69490**	**75332**	**92332**

- Test OPS in different test service models based on different instance classes and the preset 100 million rows of data. For details, see the content in bold in **Table 2-13**.

 Copyright © Huawei Technologies Co., Ltd.

Table 2-13 Test data

Weak consistency: Test OPS of the database with 100 million rows of data.					
Instance Class		4 vCPUs | 16 GB	8 vCPUs | 32 GB	16 vCPUs | 64 GB	32 vCPUs | 128 GB
No.	s6	**15637**	**17452**	**27827**	**49384**
	s5	**11541**	**14272**	**17733**	**23160**
	s4	**14926**	**25745**	**33258**	**48924**
	s3	**15431**	**27476**	**32161**	**50332**
	s2	**9004**	**14565**	**19961**	**26589**
	s1	**18448**	**27976**	**33296**	**51920**

 Copyright © Huawei Technologies Co., Ltd.