Alibaba Cloud

ApsaraDB for Cassandra Performance whitepaper

Document Version: 20211231

C-J Alibaba Cloud

Legal disclaimer

Alibaba Cloud reminds you to carefully read and fully understand the terms and conditions of this legal disclaimer before you read or use this document. If you have read or used this document, it shall be deemed as your total acceptance of this legal disclaimer.

- 1. You shall download and obtain this document from the Alibaba Cloud website or other Alibaba Cloudauthorized channels, and use this document for your own legal business activities only. The content of this document is considered confidential information of Alibaba Cloud. You shall strictly abide by the confidentiality obligations. No part of this document shall be disclosed or provided to any third party for use without the prior written consent of Alibaba Cloud.
- 2. No part of this document shall be excerpted, translated, reproduced, transmitted, or disseminated by any organization, company or individual in any form or by any means without the prior written consent of Alibaba Cloud.
- 3. The content of this document may be changed because of product version upgrade, adjustment, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and an updated version of this document will be released through Alibaba Cloud-authorized channels from time to time. You should pay attention to the version changes of this document as they occur and download and obtain the most up-to-date version of this document from Alibaba Cloud-authorized channels.
- 4. This document serves only as a reference guide for your use of Alibaba Cloud products and services. Alibaba Cloud provides this document based on the "status quo", "being defective", and "existing functions" of its products and services. Alibaba Cloud makes every effort to provide relevant operational guidance based on existing technologies. However, Alibaba Cloud hereby makes a clear statement that it in no way guarantees the accuracy, integrity, applicability, and reliability of the content of this document, either explicitly or implicitly. Alibaba Cloud shall not take legal responsibility for any errors or lost profits incurred by any organization, company, or individual arising from download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, take responsibility for any indirect, consequential, punitive, contingent, special, or punitive damages, including lost profits arising from the use or trust in this document (even if Alibaba Cloud has been notified of the possibility of such a loss).
- 5. By law, all the contents in Alibaba Cloud documents, including but not limited to pictures, architecture design, page layout, and text description, are intellectual property of Alibaba Cloud and/or its affiliates. This intellectual property includes, but is not limited to, trademark rights, patent rights, copyrights, and trade secrets. No part of this document shall be used, modified, reproduced, publicly transmitted, changed, disseminated, distributed, or published without the prior written consent of Alibaba Cloud and/or its affiliates. The names owned by Alibaba Cloud shall not be used, published, or reproduced for marketing, advertising, promotion, or other purposes without the prior written consent of Alibaba Cloud. The names owned by Alibaba Cloud and/or its affiliates Cloud include, but are not limited to, "Alibaba Cloud", "Aliyun", "HiChina", and other brands of Alibaba Cloud and/or its affiliates, which appear separately or in combination, as well as the auxiliary signs and patterns of the preceding brands, or anything similar to the company names, trade names, trademarks, product or service names, domain names, patterns, logos, marks, signs, or special descriptions that third parties identify as Alibaba Cloud and/or its affiliates.
- 6. Please directly contact Alibaba Cloud for any errors of this document.

Document conventions

Style	Description	Example				
<u>↑</u> Danger	A danger notice indicates a situation that will cause major system changes, faults, physical injuries, and other adverse results.	Danger: Resetting will result in the loss of user configuration data.				
O Warning	A warning notice indicates a situation that may cause major system changes, faults, physical injuries, and other adverse results.	Warning: Restarting will cause business interruption. About 10 minutes are required to restart an instance.				
C) Notice	A caution notice indicates warning information, supplementary instructions, and other content that the user must understand.	Notice: If the weight is set to 0, the server no longer receives new requests.				
? Note	A note indicates supplemental instructions, best practices, tips, and other content.	Onte: You can use Ctrl + A to select all files.				
>	Closing angle brackets are used to indicate a multi-level menu cascade.	Click Settings> Network> Set network type.				
Bold	Bold formatting is used for buttons , menus, page names, and other UI elements.	Click OK.				
Courier font	Courier font is used for commands	Run the cd /d C:/window command to enter the Windows system folder.				
Italic	Italic formatting is used for parameters and variables.	bae log listinstanceid Instance_ID				
[] or [a b]	This format is used for an optional value, where only one item can be selected.	ipconfig [-all -t]				
{} or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}				

Table of Contents

1.Performance overview		05
------------------------	--	----

1.Performance overview

Background information

This topic describes how to perform a benchmark test for ApsaraDB for Cassandra and provides sample results based on the test. The test results may not show the optimal performance of ApsaraDB for Cassandra because the test results vary based on the kernel and cloud environments. If you want to assess the size of the ApsaraDB for Cassandra instance that best suits your business, you can use the tests that are described in this topic. The best method to assess the instance size is to run simulated workloads on an ApsaraDB for Cassandra instance. The result is more accurate than the result provided by an external benchmarking tool.

Test tool

Perform a benchmark test for ApsaraDB for Cassandra by using Yahoo Cloud Serving Benchmark (YCSB) 0.15.0. YCSB is a standard benchmarking tool. For more information, visit https://github.com/brianfrankcooper/YCSB/tree/0.15.0/cassandra.

Test environment

Purchase an ApsaraDB for Cassandra instance for testing.

Network: virtual private cloud (VPC). You must deploy the client and server in the same region and zone. Instance architecture: one cloud data center that consists of three nodes. Instance storage: a standard SSD of 400 GB for each node. The storage capacity affects the instance performance. Client for stress testing: ecs.c6.2xlarge (8 vCPUs, 16 GB). Instance specifications: all specifications that are supported by ApsaraDB for Cassandra.

Workload description

The throughput and latency of ApsaraDB for Cassandra vary based on different workloads, such as the number of fields per row and the data size of each row. In this example, the default workloada of YCSB is used for testing. You can modify YCSB parameters based on your business requirements. Keep the default values for most parameters used to test ApsaraDB for Cassandra. For more information, visit https://github.com/brianfrankcooper/YCSB/tree/0.15.0/cassandra.

Key parameters

- 10 fields per row (default).
- 1 KB per row (default).
- Read/write operation ratio: 95:5.
- Read/write consistency level: ONE (default).
- Number of replicas: specify two replicas because disks are used.
- Stress testing threads: modify based on the instance specifications. For more information, see the test results.
- recordcount: the number of rows imported. Modify the parameter based on the specifications. For more information, see the test results.
- operationcount: the number of times of stress tests. The value of this parameter is the same as that of the recordcount parameter.

The consistency level setting may affect the performance. Specify a consistency level based on your business requirements.

Procedure

Step 1. Create a test table

```
# Replace cn-shanghai-g with the cloud data center ID of the instance that you purchased. Y
ou can view the Data Center Name parameter in the ApsaraDB for Cassandra console.
create keyspace ycsb WITH replication = {'class': 'NetworkTopologyStrategy', 'cn-shanghai-g
': 2};
create table ycsb.usertable (y_id varchar primary key, field0 varchar, field1 varchar, fiel
d2 varchar, field3 varchar, field4 varchar, field5 varchar, field6 varchar, field7 varchar,
field8 varchar, field9 varchar);
```

Step 2. Install the benchmarking tool

```
wget https://github.com/brianfrankcooper/YCSB/releases/download/0.15.0/ycsb-cassandra-bindi
ng-0.15.0.tar.gz
tar -zxf ycsb-cassandra-binding-0.15.0.tar.gz
```

Step 3. Modify the workloads/workloada code

Add the following content:

```
hosts=cds-xxxxxxx-core-003.cassandra.rds.aliyuncs.com #The endpoint of the database. You c
an view the endpoint in the ApsaraDB for Cassandra console.
cassandra.username=cassandra #The account must be granted the permissions to read and write
the ycsb keyspace.
cassandra.password=123456 #If you forget the password, you can change the password in the c
onsole.
```

Step 4. Prepare data (write test)

```
nohup ./bin/ycsb load cassandra2-cql -threads THREAD_COUNT -P workloads/workloada -s > $LO G_FILE 2>&1 &
```

The result of this test shows the maximum write throughput. To test the maximum throughput, you must increase the value of \$THREAD_COUNT and check whether the throughput increases. The specifications of the client for stress testing must be medium or high.

Step 5. Perform stress testing (read and write test)

```
nohup ./bin/ycsb run cassandra2-cql -threads THREAD_COUNT -P workloads/workloada -s > $LOG FILE 2>&1 &
```

The result of this test shows the read and write performance.

Test result

The test results in this example are provided only for reference. The throughput and latency vary based on different workloads. You can use different parameters and workloads, or increase the volumes of test data (longer duration) to obtain test results that best suit your business. The client specifications may affect the test results. Do not use shared instances.

Test result description

- Load: data preparation (write test)
- Run: stress testing (read and write test)
- OPS: operations per second, which indicates the overall throughput.
- WAVG: the average write latency. Unit: microseconds.
- RAVG: the average read latency. Unit: microseconds.
- RP999: the write latency of the 99.9th percentile. Unit: microseconds.
- Thread: 100/100, which indicates the number of testing threads in the data preparation phase compared with the number of testing threads in the stress testing phase.

A full load test and a regular test are performed during the stress testing phase.

80% CPU load

Spec ifica tion	T hre ad	Data volu me (ten thou sand rows)	Load OPS	Load WAV G	Run OPS	Run WAV G	Run RAV G	Run RP95	Run RP99	Run RP99 9
4 vCPUs 8 GB	100/1 00	1600	32277	3071	29745	2846	3363	7795	23039	43999

60% CPU load

Spec ifica tion	T hre ad	Data volu me (ten thou sand rows)	Load OPS	Load WAV G	Run OPS	Run WAV G	Run RAV G	Run RP95	Run RP99	Run RP99 9
4 vCPUs 8 GB	100/1 6	1600	32063	3093	16721	514	974	1879	3047	28063

? Note

This topic lists the test results of an instance that uses standard SSDs. Ultra disks also provide high IOPS. When the amount of data and the instance specifications are small, the impact of ultra disks on performance is close to that of SSDs on performance. The storage is no longer the performance bottleneck. Therefore, ultra disks are not used in the tests. You can use simulated workloads based on your business scenario to obtain more accurate results. The impacts of applications must also be taken into account. For example, the garbage collection mechanism of Java clients may increase the latency.