

Alibaba Cloud ApsaraDB for Cassandra

Performance whitepaper

Issue: 20200521









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 Cloud-authorized 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 due to product version upgrades, adjustments, or other reasons. Alibaba Cloud reserves the right to modify the content of this document without notice and the updated versions of this document will be occasionally released through Alibaba Cloud-authorized channels. You shall 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 the document in the context that Alibaba Cloud products and services are provided on an "as is", "with all faults" and "as available" basis. 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 bear any liability for any errors or financial losses incurred by any organizations, companies, or individuals arising from their download, use, or trust in this document. Alibaba Cloud shall not, under any circumstances, bear responsibility for any indirect, consequential, exemplary, incidental, 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 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 contact Alibaba Cloud directly if you discover any errors in this document.

Document conventions

Style	Description	Example
	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.
	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.
	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.
	A note indicates supplemental instructions, best practices, tips, and other content.	 Note: 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 <code>cd /d C:/window</code> command to enter the Windows system folder.
Italic	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid Instance_ID</code>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>

Style	Description	Example
{ } or {a b}	This format is used for a required value, where only one item can be selected.	switch {active stand}

Contents

Legal disclaimer.....	I
Document conventions.....	I
1 Performance white paper.....	1

1 Performance white paper

This topic describes how to perform benchmark tests for ApsaraDB for Cassandra and provides sample results. The test results may vary as the kernel and cloud environments are constantly optimized. If you need to estimate the ApsaraDB for Cassandra instance scales suitable for your business, you can refer to the test approaches described in this topic. To obtain a precise result, we recommend that you run simulated services on instances.

Test tool

Yahoo! Cloud Serving Benchmark (YCSB) 0.15.0 (new released). For more information, visit <https://github.com/brianfrankcooper/YCSB/tree/0.15.0/cassandra>.

Test environment

Buy an ApsaraDB for Cassandra instance of the following specifications:

Network: VPC. Deploy the client and server in the same region and zone. Instance architecture: a data center that contains three nodes. Instance storage: 400 GB per node and standard SSD. Stress testing server type: ecs.c6.2xlarge (8 vCPUs, 16 GB). Instance specifications: all specifications supported by ApsaraDB for Cassandra.

Test loads

Different services bear different loads such as the number of fields and data volume in each row and result in different throughput and latency. This topic uses the default [workload](#) provided by YCSB for testing. You can customize YCSB parameters based on your service. Most ApsaraDB for Cassandra parameters use the default values. For more information, visit <https://github.com/brianfrankcooper/YCSB/tree/0.15.0/cassandra>.

Parameters

- 10 fields per row (default).
- 1 KB records per row (default).
- Read/write operation ratio: 95:5.
- Read/write consistency level: ONE (default).
- Number of replicas: Configure two replicas because the standard SSD is used.
- Stress testing threads: modified based on the instance specifications. For more information, see the test results.

- recordcount: the row number of imported data. It is modified based on the specifications. For more information, see the test results.
- operationcount: the number of stress testing operations. It is the same as recordcount. <https://github.com/brianfrankcooper/YCSB/tree/0.15.0/cassandra>

Note that adjusting the consistency level will affect the performance. Specify a consistency level based on your business requirements.

Test procedure

1. Create a test table

```
# Replace cn-shanghai-g with the data center ID of the tested instance. You 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,
field2 varchar, field3 varchar, field4 varchar, field5 varchar, field6 varchar, field7 varchar
, field8 varchar, field9 varchar);
```

2. Install the test tool

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

3. Edit the workloads/workloada code

Add the following lines of code to workloada:

```
hosts=cds-xxxxxxx-core-003.cassandra.rds.aliyuncs.com #The endpoint of the instance
. You can view it in the ApsaraDB for Cassandra console.
cassandra.username=cassandra #The account of the instance. The account must have
permissions to read and write ycsb keyspace.
cassandra.password=123456 #The password of the account. You can change the
password in the console.
```

4. Prepare data (read performance test)

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

You can view the maximum throughput based on the test result. To test the maximum throughput, you must increase the value of \$THREAD_COUNT and view whether the throughput increases. We recommend that you select medium or large specifications for the stress testing server.

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

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

You can view the read and write performance based on the test result.

Test results

The test results are for reference only. Different loads will result in different throughput and latency. You can use various parameters, loads, data volumes, and time periods to obtain test results that suit your business. The client specifications will affect the performance. Do not use shared instances.

Result description

Load: the data preparation phase (read performance test). Run: the stress testing phase (read and write performance test). OPS: the operations per second, indicating the throughput in all phases. WAVG: the average write latency. Unit: microseconds. RAVG: the average read latency. Unit: microseconds. RP999: the 99.9% read latency. Unit: microseconds. Thread: the number of YCSB testing threads during data preparation/the number of YCSB testing threads during stress testing. For example, 100/100.

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

80% CPU load

Specification	Thread	Data volume (ten thousand rows)	Load OPS	Load WAVG	Run OPS	Run WAVG	Run RAVG	Run RP95	Run RP99	Run RP999
4 cores, 8 GB	100/100	1,600	32,277	3,071	29,745	2,846	3,363	7,795	23,039	43,999

60% CPU load

Specification	Thread	Data volume (ten thousand rows)	Load OPS	Load WAVG	Run OPS	Run WAVG	Run RAVG	Run RP95	Run RP99	Run RP999
4 cores, 8 GB	100/16	1,600	32,063	3,093	16,721	514	974	1,879	3,047	28,063

Note: This topic lists only the test results of an instance based on standard SSD. The ultra disk provides high IOPS and brings little impact on performance when you test instances of the basic specifications by using a small amount of data. Therefore, instances based on ultra disks are not provided for reference. You can simulate loads based on your business during stress tests to obtain precise results. The impacts of applications must also be considered. For example, the garbage collection mechanism of Java clients causes latency.