

Alibaba Cloud

ApsaraDB for MongoDB Performance White Paper

Document Version: 20210303

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 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 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
 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.
 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.
 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.	 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.
<i>Italic</i>	Italic formatting is used for parameters and variables.	<code>bae log list --instanceid</code> <i>Instance_ID</i>
[] or [a b]	This format is used for an optional value, where only one item can be selected.	<code>ipconfig [-all -t]</code>
{ } or {a b}	This format is used for a required value, where only one item can be selected.	<code>switch {active stand}</code>

Table of Contents

1.Overview	05
2.Test environment	06
3.Test tool	07
4.Test procedure	08
5.Test results	09

1. Overview

ApsaraDB for MongoDB fully supports the MongoDB protocol and is based on the distributed Apsara system and high-reliability storage engine. ApsaraDB for MongoDB provides solutions such as multiple-node high-availability architecture, elastic scaling, disaster recovery, backup and rollback, and performance optimization.

For more information, see [What is ApsaraDB for MongoDB?](#).


2. Test environment

This topic describes the performance test environment for ApsaraDB for MongoDB.

- All tests are performed in Zone B of the China (Hangzhou) region.
- A replica set instance (three-node) is used.
- The database version of the instance is 3.4.
- The network type of the instance is classic network.
- The CentOS 6.0 64-bit image is used in the stress testing.

3. Test tool

The open-source YCSB 0.12.0 tool is used for stress testing.


 **Note** YCSB is a Java tool that can be used in performance testing for multiple databases. For more information about its installation and usage, see [YCSB](#).

4. Test procedure

This topic describes how to test the performance of ApsaraDB for MongoDB.

Procedure

1. Modify the following parameters in the workload configuration: recordcount, operationcount, readproportion, and updateproportion.

 **Note** The values of the recordcount and operationcount parameters vary with instance types. For more information, see [Test results](#).

2. Prepare data resources for testing.

```
./bin/ycsb load mongodb -s -P workloads/workloada -p mongodb.url=mongodb://ip:port/ycsb? w=0 -threads xx > outputLoad.txt
```

3. Start the performance test.


```
./bin/ycsb run mongodb -s -P workloads/workloada -p mongodb.url=mongodb://ip:port/ycsb? w=0 -threads xx > outputRun.txt
```


5. Test results

This topic describes the performance test results of ApsaraDB for MongoDB instances at different read/write ratios.

Metrics

- **Count:** the sum of the recordcount value and the operationcount value. The recordcount parameter indicates the number of existing records. The operationcount parameter indicates the number of the operations to be executed.
- **Threads:** the number of threads that are used for the client test.

 **Note** Four Elastic Compute Service (ECS) instances of 4 cores and 32 GB memory are used to evenly distribute threads for client concurrent testing.

- **Throughput:** the number of read and write operations. Unit: ops/s.
- **RAL:** the average latency of read operations. Unit: us.
- **WAL:** the average latency of write operations. Unit: us.

Test results at 50:50 read/write ratio

Instance type	Count	Thread	Throughput	RAL	WAL
General-purpose, 1 core and 2 GB memory	1,000,000	100	3,997	22,080	27,934
General-purpose, 2 cores and 4 GB memory	2,000,000	100	7,674	11,778	14,271
General-purpose, 4 cores and 8 GB memory	4,000,000	100	17,002	5,249	6,502
General-purpose, 8 cores and 16 GB memory	8,000,000	100	30,500	3,027	3,520
General-purpose, 8 cores and 32 GB memory	16,000,000	100	33,655	2,679	3,253
General-purpose, 16 cores and 64 GB memory	32,000,000	100	64,883	1,322	1,761
Dedicated, 2 cores and 16 GB memory	100,000,000	150	4,354	30,674	38,167
Dedicated, 4 cores and 32 GB memory	100000000	150	10,890	12,517	15,019

Instance type	Count	Thread	Throughput	RAL	WAL
Dedicated, 8 cores and 64 GB memory	100,000,000	150	21,145	6,347	7,826
Dedicated, 16 cores and 128 GB memory	100,000,000	150	50,625	2,589	3,323
Dedicated, 32 cores and 256 GB memory	100,000,000	150	65,472	1,982	2,588
Dedicated host, 30 cores and 220 GB memory	100,000,000	150	62,472	1,955	2,770
Dedicated host, 60 cores and 440 GB memory	100,000,000	150	90,181	1,410	1,870

Test results at 95:5 read/write ratio

Instance type	Count	Thread	Throughput	RAL	WAL
General-purpose, 1 core and 2 GB memory	1,000,000	100	7,849	12,519	16,801
General-purpose, 2 cores and 4 GB memory	2,000,000	100	14,923	6,621	8,109
General-purpose, 4 cores and 8 GB memory	4,000,000	100	37,573	2,623	3,277
General-purpose, 8 cores and 16 GB memory	8,000,000	100	51,085	1,936	2,247
General-purpose, 8 cores and 32 GB memory	16,000,000	100	70,780	1,383	1,885
General-purpose, 16 cores and 64 GB memory	32,000,000	100	105,606	920	1,371
Dedicated, 2 cores and 16 GB memory	100,000,000	150	7,175	20,701	24,635
Dedicated, 4 cores and 32 GB memory	100,000,000	150	17,270	8,634	9,529

Instance type	Count	Thread	Throughput	RAL	WAL
Dedicated, 8 cores and 64 GB memory	100,000,000	150	46,707	3,167	3,920
Dedicated, 16 cores and 128 GB memory	100,000,000	150	106,386	1,372	2,013
Dedicated, 32 cores and 256 GB memory	100,000,000	150	150,378	970	1,233
Dedicated host, 30 cores and 220 GB memory	100,000,000	150	132,717	1,100	1,405
Dedicated host, 60 cores and 440 GB memory	100,000,000	150	225,365	643	856