Alibaba Cloud ApsaraDB for POLARDB

Performance White Paper

Issue: 20190819

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 content of the Alibaba Cloud website, including but not limited to works, products, images, archives, information, materials, website architecture, website graphic layout, and webpage design, 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 the Alibaba Cloud website, product programs, or content 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.

Generic conventions

Table -1: Style conventions

Style	Description	Example
	This warning information 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.
	This warning information 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 restore business.
	This indicates warning information, supplementary instructions, and other content that the user must understand.	 Notice: Take the necessary precautions to save exported data containing sensitive information.
	This indicates supplemental instructions, best practices, tips, and other content that is good to know for the user.	 Note: You can use Ctrl + A to select all files.
>	Multi-level menu cascade.	Settings > Network > Set network type
Bold	It is used for buttons, menus , page names, and other UI elements.	Click OK.
`Courier font`	It is used for commands.	Run the `cd / d C :/ windows` command to enter the Windows system folder.
Italics	It is used for parameters and variables.	`bae log list --instanceid` *Instance_ID*
[] or [a|b]	It indicates that it is a optional value, and only one item can be selected.	`ipconfig` *[-all|-t]*

Style	Description	Example
{} or {a|b}	It indicates that it is a required value, and only one item can be selected.	`swich {stand | slave}`

Contents

1 Performance white paper

This topic describes how to evaluate the optimal performance of POLARDB using Sysbench 0.5.

Sysbench is a modular, cross-platform and multi-threaded benchmark tool for evaluating important OS parameters of a system running a database under intensive load. The benchmark tool is designed to quickly evaluate the database system performance without complex database benchmark settings or even without database installations.

Prerequisites

- Three ECS instances:
  - Each instance has 32 vCPUs. For example, you can use ecs.sn1ne.8xlarge instances.
  - The ECS instances and POLARDB cluster must be in the same region and zone.
  - You must set usernames and passwords for the ECS instances.
  - The OS must be CentOS 7.4 (64-bit).
- A POLARDB cluster:
  - The cluster contains one primary instance and one read-only instance. You must add more read-only instances if you need to evaluate the performance of a cluster with multiple read-only instances.
  - You must set a username and a password for the cluster.
  - You must add internal IP addresses of the ECS instances to the whitelist of your POLARDB cluster.

Install Sysbench 0.5

1. To install Sysbench 0.5, run the following commands in ECS.

```
yum install gcc gcc - c ++ autoconf automake make
libtool bzr mysql - devel git mysql

git clone https :// github . com / akopytov / sysbench . git

cd sysbench

git checkout 0 . 5

./ autogen . sh

./ configure -- prefix =/ usr -- mandir =/ usr / share / man

make

make install
```

2. Run the following commands to configure the Sysbench client so that it uses all CPU cores available to process the data (two CPU cores are used by default), which will reduce cross-core context switching.

```
sudo sh - c ' for x in / sys / class / net / eth0 / queues /
rx -*; do echo ffffffff >$ x / rps_cpus ; done '
```

Note: ffffffff indicates 32 cores are used for data processing. Modify the command to reflect your actual configurations. For example, enter ff if your ECS instance has eight cores.

```
sudo sh - c " echo 32768 > / proc / sys / net / core /
rps_sock_f low_entrie s "

sudo sh - c " echo 4096 > / sys / class / net / eth0 / queues
/ rx - 0 / rps_flow_c nt "

sudo sh - c " echo 4096 > / sys / class / net / eth0 / queues
/ rx - 1 / rps_flow_c nt "
```

Test procedure

1. Obtain the cluster connection string and port number.
   a. Log on to the POLARDB console, and enter the Clusters page.
   b. Click the cluster ID, or click Manage to enter the Cluster Information page.
   c. Find the ../DNPOLA1840729/EN-US_TP_3018.dita#concept_imd_wlq_tdb and port number.
2. Run the following commands in ECS to create a sbtest database in the POLARDB instance.

```
mysql -h XXX -P XXX -u XXX -p XXX -e 'create database sbtest'
```

> **Note:**
> Replace XXX in the command and subsequent commands with your POLARDB cluster connection string (VPC), port number, username, and password, respectively.

3. Preparing test data: You can create a table in the database using Sysbench, and insert data into the table.

```
sysbench --test=sysbench/tests/db/oltp.lua --mysql-host=XXX --mysql-port=XXX --mysql-user=XXX --mysql-password=XXX --mysql-db=sbtest --mysql-table-engine=innodb --oltp-table-size=25000 --oltp-tables-count=250 --db-driver=mysql prepare
```

4. Evaluate the read performance of the database using Sysbench. The evaluation will take 10 minutes.

```
sysbench --test=sysbench/tests/db/oltp.lua --mysql-host=XXX --oltp-tables-count=250 --mysql-user=XXX --mysql-password=XXX --mysql-port=XXX --db-driver=mysql --oltp-tablesize=25000 --mysql-db=sbtest --max-requests=0 --oltp_simple_ranges=0 --oltp-distinct-ranges=0 --oltp-sum-ranges=0 --oltp-order-ranges=0 --max-time=600 --oltp-read-only=on --num-threads=500 run
```

5. Evaluate the write performance of the database using Sysbench. The evaluation will take 10 minutes.

```
sysbench --test=sysbench/tests/db/oltp.lua --mysql-host=XXX --oltp-tables-count=250 --mysql-user=XXX --mysql-password=XXX --mysql-port=XXX --db-driver=mysql --oltp-tablesize=25000 --mysql-db=sbtest --max-requests=0 --max-time=600 --oltp_simple_ranges=0 --oltp-distinct-ranges=0 --oltp-sum-ranges=0
```

```
-- oltporder - ranges = 0  -- oltp - point - selects = 0  -- num -
threads = 128  -- randtype = uniform   run
```

6. During the tests, you can connect to the ECS instance in a new window, and run htop command to check if the CPU utilization of the Sysbench client is normal.

```
yum   install   htop

htop
```

Note:

- After running the htop command, you can click Q to exit.

- For more information about htop, go to http://hisham.hm/htop/?spm=a2c4g.11186623.2.6.eKuBNC.

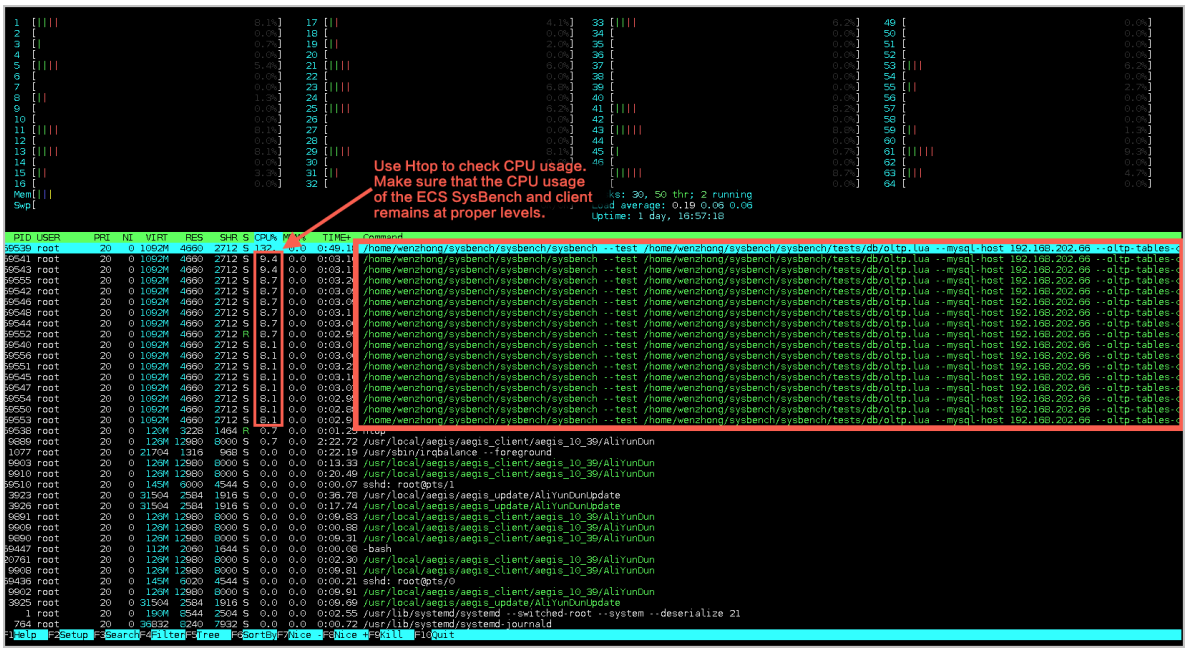

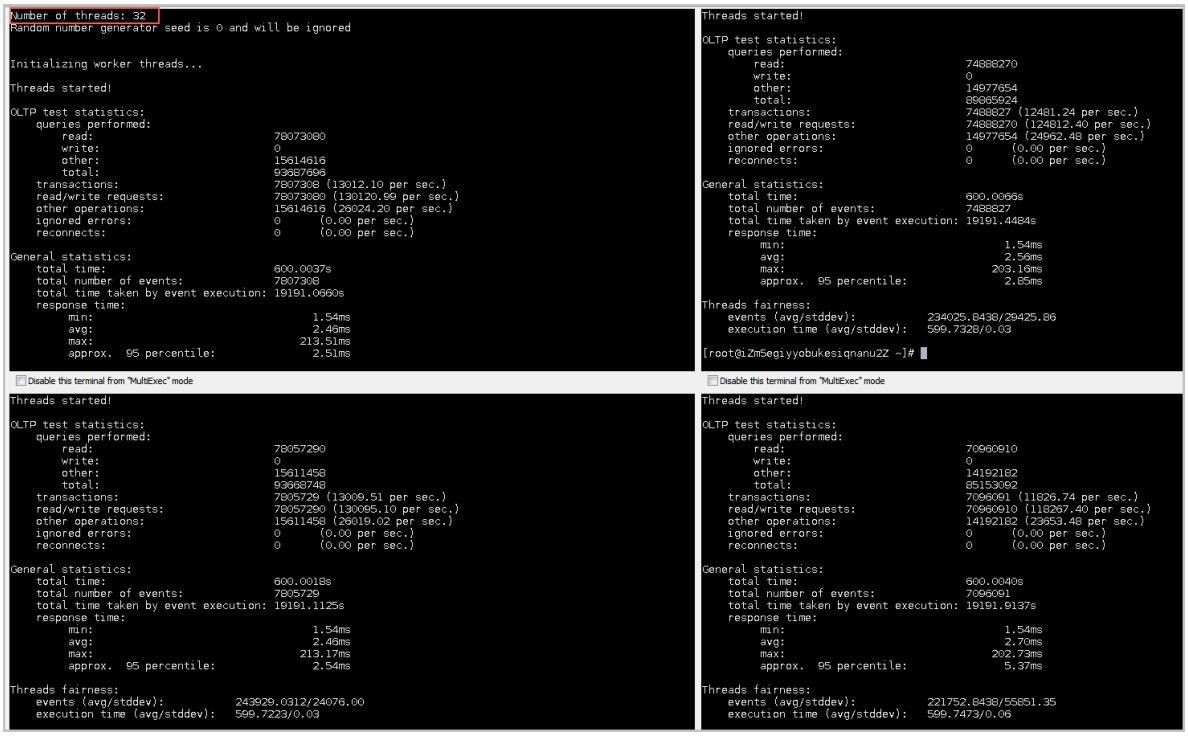

Test results

Obtain the QPS and TPS from the test results in the log file.

Note:

QPS indicates the number of SQL statements (including INSERT, SELECT, UPDATE, and DELETE) executed by a database every second.

Test results for a single ECS instance:

```
sysbench 0.5:  multi-threaded system evaluation benchmark

Running the test with following options:
Number of threads: 128
Random number generator seed is 0 and will be ignored


Initializing worker threads...

Threads started!

OLTP test statistics:
    queries performed:
        read:                            313320510  --Number of read operations
        write:                           0          --Number of write operations
        other:                           62664102   --Numbe of other operations such as COMMIT excluding CURD
        total:                           375984612  --Total number of operations
    transactions:                        31332051   (52219.88 per sec.)  -- Total number of transactions per second
    read/write requests:                 313320510  (522198.79 per sec.)  Number of read and write operations per second
    other operations:                    62664102   (104439.76 per sec.)  -- Number of other operations per second
    ignored errors:                      0          (0.00 per sec.)
    reconnects:                          0          (0.00 per sec.)

General statistics:                      --| The database performance is measured based on the QPS, namely, the number of read and write
    total time:                          600.0024s  -- Total time taken                                  operations per second.
    total number of events:              31332051   -- Total number of transactions
    total time taken by event execution: 76751.6764s --Total time taken by all transactions (with concurrency not considered)
    response time:
         min:                                  1.71ms -- Minimum time taken
         avg:                                  2.45ms -- Average time taken
         max:                                194.31ms -- Maximum time taken
         approx.  95 percentile:               2.67ms -- Average time taken by 95% of the transactions processed

Threads fairness:
    events (avg/stddev):           244781.6484/25277.49 -- Total number of transactions processed/Standard offset
    execution time (avg/stddev):   599.6225/0.14 --Total time taken/Standard offset
```

Total QPS for the three ECS instances:

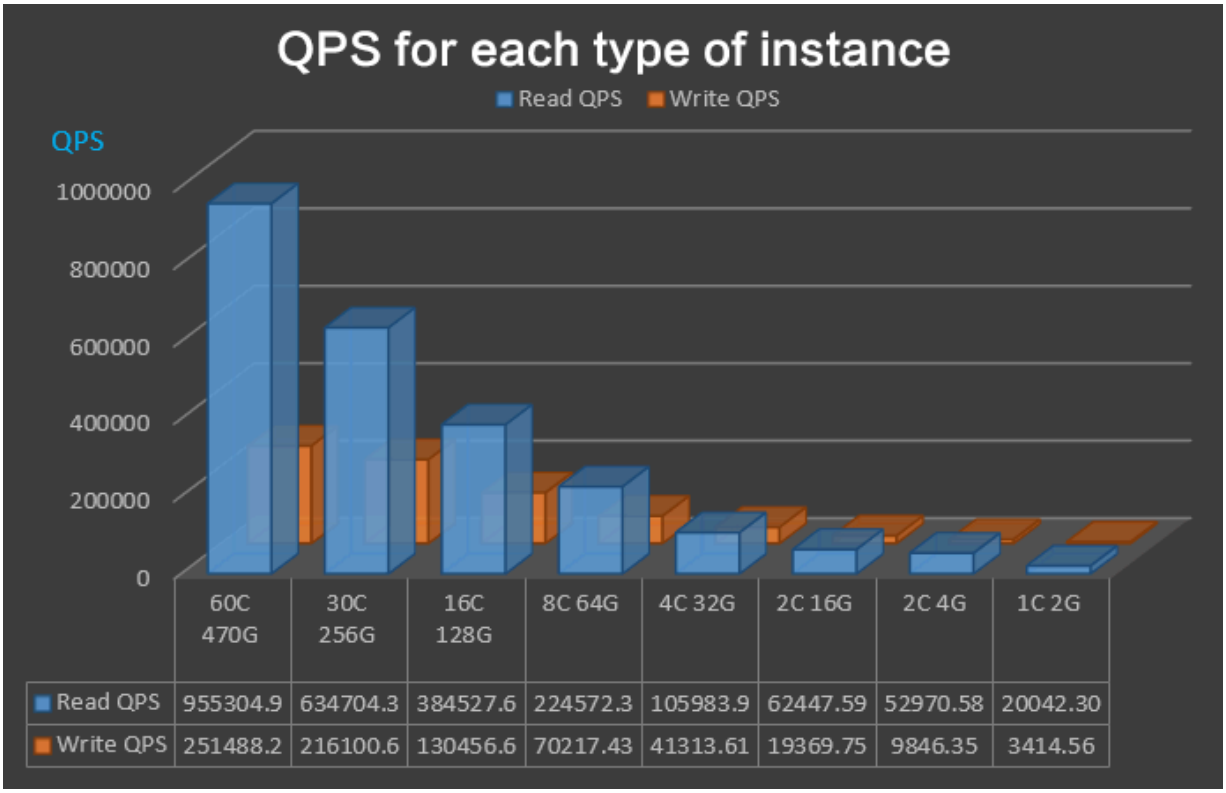

Total TPS for the three ECS instances:

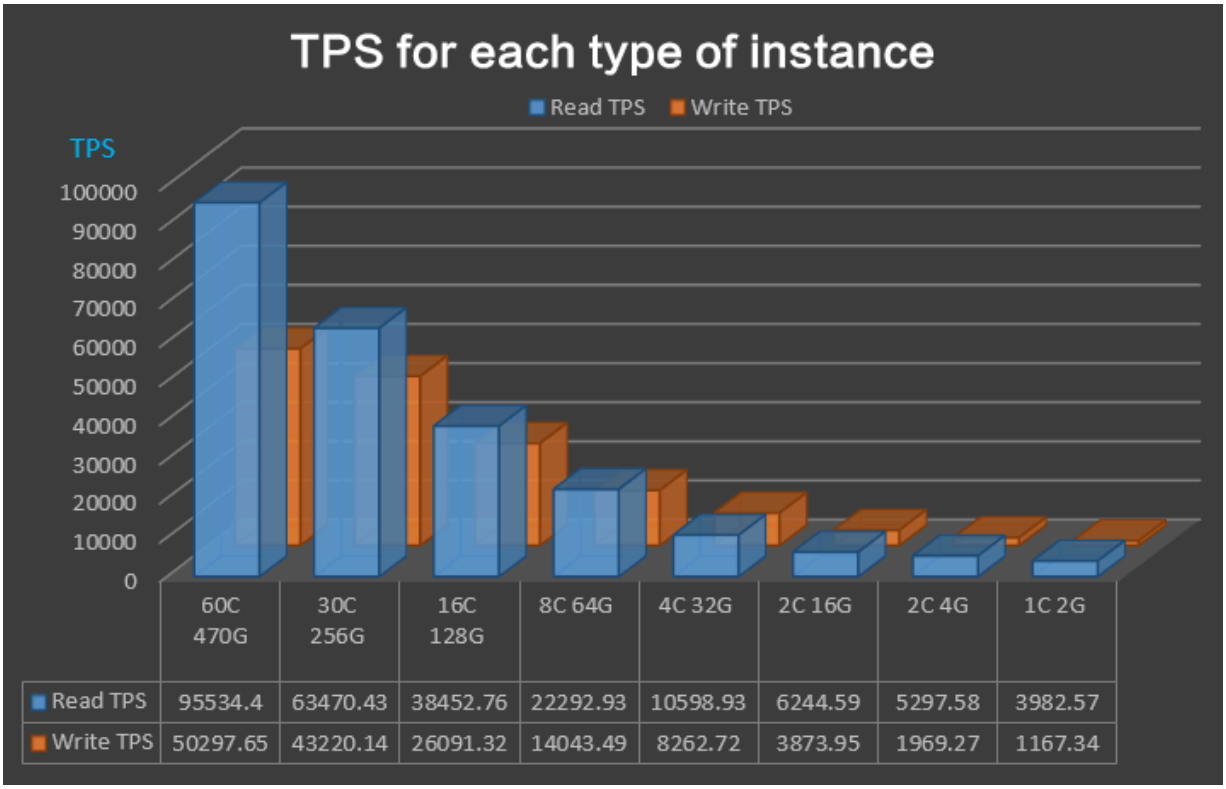

	60C 470G	30C 256G	16C 128G	8C 64G	4C 32G	2C 16G	2C 4G	1C 2G
Read TPS	95534.4	63470.43	38452.76	22292.93	10598.93	6244.59	5297.58	3982.57
Write TPS	50297.65	43220.14	26091.32	14043.49	8262.72	3873.95	1969.27	1167.34

The following figure shows the QPS for multiple read-only instances. Five read-only instances are used in this example, each with 4 cores and 32 GB RAM.

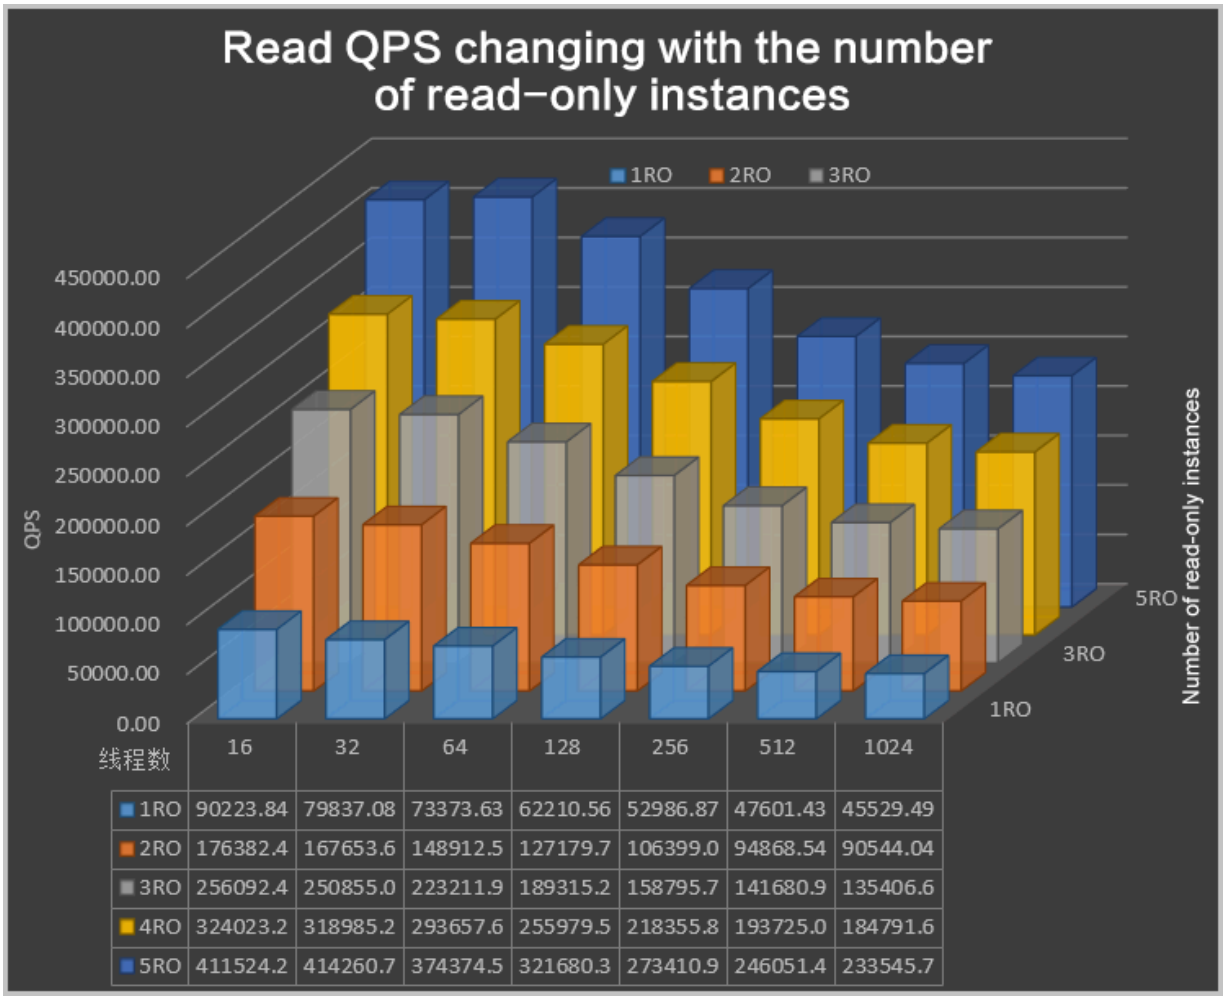

线程数	16	32	64	128	256	512	1024
1RO	90223.84	79837.08	73373.63	62210.56	52986.87	47601.43	45529.49
2RO	176382.4	167653.6	148912.5	127179.7	106399.0	94868.54	90544.04
3RO	256092.4	250855.0	223211.9	189315.2	158795.7	141680.9	135406.6
4RO	324023.2	318985.2	293657.6	255979.5	218355.8	193725.0	184791.6
5RO	411524.2	414260.7	374374.5	321680.3	273410.9	246051.4	233545.7