

iSCSI Performance Tuning

Best Practice Guide

Version 1.0

January 02, 2020

Table of Content

| | |
|---|----|
| TABLE OF CONTENT..... | 1 |
| ANNOUNCEMENT | 2 |
| Edition 1.0 (JAN. 2020) | 2 |
| Trademarks | 2 |
| NOTICES..... | 3 |
| PREFACE..... | 4 |
| About This Guide | 4 |
| Technical Support..... | 4 |
| Information, Tip and Caution..... | 5 |
| 1. OVERVIEW..... | 6 |
| 1.1 Audience | 6 |
| 2. STORAGE BEST CONFIGURATION..... | 6 |
| 2.1. Network Configuration | 6 |
| 2.2. RAID Configuration | 7 |
| 3. SERVER CONFIGURATION | 9 |
| 3.1. Network Configuration | 9 |
| 3.2. Local Drive..... | 9 |
| 3.3. Windows Server OS | 9 |
| 3.4. Unix-Like OS | 11 |
| 4. SUPPORT AND OTHER RESOURCES | 13 |
| 4.1. Getting Technical Support | 13 |
| 4.2. Documentation Feedback..... | 14 |
| APPENDIX..... | 15 |
| End-User License Agreement (EULA) | 15 |

Announcement

© Copyright 2020 QSAN Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from QSAN Technology, Inc.

Edition 1.0 (JAN. 2020)

This edition applies to QSAN XCubeSAN SAN storage systems. QSAN believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Trademarks

QSAN, the QSAN logo, XCubeSAN, and QSAN.com are trademarks or registered trademarks of QSAN Technology, Inc.

Intel, Xeon, Pentium, Core, Celeron, Intel Inside, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Other trademarks and trade names used in this document to refer to either the entities claiming the marks and names, or their products are the property of their respective owners.

Notices

The information contained in this manual has been reviewed for accuracy. But it could include typographical errors or technical inaccuracies. Changes are made to the document periodically. These changes will be incorporated in new editions of the publication. QSAN may make improvements or changes in the products. All features, functionality, and product specifications are subject to change without prior notice or obligation. All statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products.

All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

Preface

About This Guide

This performance tuning guide provides technical guidance for adjusting/increasing the performance when testing QSAN products which provides iSCSI service, and it is intended for use by system administrators, SAN/NAS designers, storage consultants, or anyone who has purchased these products and is familiar with server and computer network, network administration, storage system installation and configuration, network attached storage management, and relevant protocols.



CAUTION:

Do NOT attempt to service, change, disassemble or upgrade the equipment's components by yourself. Doing so may violate your warranty and expose you to electric shock. Refer all servicing to authorized service personnel. Please always follow the instructions in this owner's manual.

Technical Support

Do you have any questions or need help troubleshooting a problem? Please contact QSAN Support, we will reply to you as soon as possible.

- Via the Web: http://www.qsan.com/en/contact_support.php
- Via Telephone: +886-2-7720-6355
- (Service hours: 09:30 - 18:00, Monday - Friday, UTC+8)
- Via Skype Chat, Skype ID: qsan.support
- (Service hours: 09:30 - 02:00, Monday - Friday, UTC+8, Summertime: 09:30 - 01:00)
- Via Email: support@qsan.com

Information, Tip and Caution

This manual uses the following symbols to draw attention to important safety and operational information.



INFORMATION:

INFORMATION provides useful knowledge, definition, or terminology for reference.



TIP:

TIP provides helpful suggestions for performing tasks more effectively.



CAUTION:

CAUTION indicates that failure to take a specified action could result in damage to the system.

1. Overview

1.1 Audience

This document is applicable for those technical members who are familiar with QSAN products, and for those who are good at network trouble shooting, Windows Server and Unix-Like OS operations, and basic hardware installations.

Please read this document carefully before trying to adjust any parameter on server side. Doing an adjustment with wrong understanding may lead you to get a worse performance experience than ever. If you have any questions about the adjustment, please consult QSAN Technical Support for further assistance.

2. Storage Best Configuration

2.1. Network Configuration

Please make sure the followings are configured properly based on user's environment, you don't need to set each setting to be exactly the same, please try to adjust and find the best ones for user:

- Make certain that each network port has been used for I/O transferring, you may check from Performance Monitor page on web UI for reference
- If you have enabled Jumbo Frame from controller side, make sure all ports (including switch

- Don't change anything when creating a volume, you may need to enable Video Editing Mode if user's environment needs a more stable performance.

Create Volume

General
Advanced
Summary

Volume Advanced Settings

Please configure the volume advanced settings.

Block Size : 512 Byte

Priority : High

The priority is the comparison with the other volumes.

Background I/O Priority : High

Background I/O priority will influence volume initialization, rebuild, and migration.

Erase Volume Data : Fast Erase

☒ Enable Cache Mode (Write-back Cache)
Write back optimizes the system speed but comes with the risk which the data may be inconsistent be short time interval.

☐ Enable Video Editing Mode
Please enable it when the application is in the video editing environment. It sacrifices a bit of performance for a more stable performance.

☒ Enable Read-ahead
The system will identify what is needed next, based on the content just retrieved from the disk, and the buffer. When the data to be transmitted is continuous, this feature will improve performance.

- The overall performance would be a bit lower if Video Editing Mode or Fast Rebuild is enabled, but it won't be a problem to use if user needs it.
- (Optional) This is an optional configuration, but it is important to realize, if overall performance (theory) from disk drives exceeds below thresholds:
- Don't enable/take snapshot on the volume that you would like to verify performance, QSAN's Snapshot mechanism uses differential copy (on write), write performance will be low when testing write performance on a volume which has taken a snapshot.
- (Optional) This is an optional configuration, but it is important to realize, if overall performance (theory) from disk drives exceeds below thresholds:
 - 6400MB/s for read (Sequential)
 - 2500MB/s for write (Sequential)
 - 600000 IOPS for read (Sequential)
 - 300000 IOPS for write (Sequential)
 - 300000 IOPS for read (Random)
 - 100000 IOPS for write (Random)

Please try to divide disk drives to two parts with the same quantity, and create two RAID pools, then assign to controller1 and controller2 separately.

3. Server Configuration

3.1. Network Configuration

Nowadays, those mainstream OSes contain the driver by default for most popular HBA cards, but it is still recommended to update the driver version to the latest one for the installed HBA on the server side, in case of any incompatibility between the OS and HBA card.

- Make certain that each network port has been used for I/O transferring, you may check from Task Manager from Windows Server.
- You may specify source IP address when logging in iSCSI target from Windows Server iSCSI Initiator, to ensure that each NIC port on server side will be used.
- Unix-Like OS may not support to specify a source IP address when logging in a iSCSI target, so you may need to configure different network segments for each NIC port on server side.
- Make sure a proper network cable with correct link speed is used for each NIC port.

**TIP:**

If you have a switch connected between server and storage, make sure the Flow Control on the switch side is ON, and be sure the other settings are the same as server's NIC.

3.2. Local Drive

If you are going to test the performance by copy and paste operation, please make sure that the local drive installed on your server possesses high performance, as a copy and paste operation relates to read and write performance from local drive as well. E.g., if local drive has only 100MB/s performance for read, while the RAID volume could deliver up to 1000MB/s performance for read and write, the test result by a copy and paste will be around 100MB/s when copying data from the local drive and pasting to the RAID volume, which means the bottleneck is on the local drive instead.

3.3. Windows Server OS

After the above configurations are confirmed, you may or may not encounter slow performance when testing read/write performance with through iSCSI connections, if you were suffering with slow performance, please try to adjust below parameters on server's HBA driver:

- Visit Device Manager, in Network adapters, right-click the NIC port you are using for performance testing, choose Properties, go to Advanced, find RSS Queue (different brands may have different terms), adjust the RSS Queue to 2 from 8.
- In the same location as above, find Receive Buffers, adjust the value to be maximum.

- In the same location as above, find Transmit Buffers, adjust the value to be maximum.
- In the same location as above, find Interrupt Moderation Rate, adjust the value to be Off.

If the above adjustments do not help, please go further for the below option:

- Open a Command Line, type “netsh int tcp set global autotuninglevel=restricted” or type “netsh int tcp set global autotuninglevel=highlyrestricted”.



INFORMATION:

It is necessary to re-log in the connected iSCSI session after the above parameters are made, a reboot of server may also be a must if it seems doesn't work after a re-login of iSCSI session.



CAUTION:

Please verify the performance right after one of adjustments above is made, to find out which one is suitable for user's environment, it is not necessary to adjust each item.

3.4. Unix-Like OS

After the above configurations are confirmed, you may or may not encounter slow performance when testing read/write performance with through iSCSI connections, if you were suffering with slow performance, please try to adjust below parameters on server:

- Please firstly finish all the configurations mentioned in this [white paper](#).
- Adjust Read Ahead (RA) buffer for the RAID volume via blockdev command:

```
root@antony-virtual-machine:/home/antony# blockdev --report
RO   RA   SSZ   BSZ   StartSec   Size   Device
rw   256  2048  2048       0   1051721728 /dev/sr0
rw   256   512  4096       0   17179869184 /dev/sda
rw   256   512  4096   2048   16105078784 /dev/sda1
rw   256   512  1024  31459326   1024   /dev/sda2
rw   256   512  4096  31459328   1071644672 /dev/sda5
root@antony-virtual-machine:/home/antony# blockdev --setra 4096 /dev/sda
```



Set the RA value to be higher like 4096 or 8192 for each RAID volume, if the RAID volume is a multipath'ed device, please adjust the dm-x device.

- Adjust the Receive Buffer for TCP/IP:

```
root@antony-virtual-machine:/home/antony# sysctl -w net.core.rmem_max=524284
net.core.rmem_max = 524284
root@antony-virtual-machine:/home/antony#
```

Set the receive buffer to be 524284 or higher (twice the value).

- Disable HT (Hyper Threading in BIOS of server):





INFORMATION:

It is necessary to re-log in the connected iSCSI session after the above parameters are made.



CAUTION:

Please verify the performance right after one of adjustments above is made, to find out which one is suitable for user's environment, it is not necessary to adjust each item.

4. Support and Other Resources

4.1. Getting Technical Support

After installing your device, locate the serial number on the sticker located on the side of the chassis or register your product at partner.qsan.com/ (End-User Registration). We recommend registering your product in QSAN partner website for firmware updates, document download, and latest news in e-blast. To contact QSAN Support, please use the following information.

1. Via the Web: http://www.qsan.com/en/contact_support.php
2. Via Telephone: +886-2-7720-6355
(Service hours: 09:30 - 18:00, Monday - Friday, UTC+8)
3. Via Skype Chat, Skype ID: qsan.support
(Service hours: 09:30 - 02:00, Monday - Friday, UTC+8, Summer time: 09:30 - 01:00)
4. Via Email: support@qsan.com

Information to collect

1. Product name, model or version, and serial number
2. Firmware version
3. Error messages or screenshot images
4. Product-specific reports and logs
5. Add-on products or components installed
6. Third-party products or components installed

Information for Technical Support

The following system information is necessary for technical support; please refer to following for what and where to get the information of your XCubeSAN model.

If the technical support requests you to download the service log/debug info, please navigate on the **SANOS UI → System Settings → Maintenance → System Information**, then click on the **“Download Service Package”** button for download.

The screenshot displays the QSAN SANOS 4.0 web interface. The left sidebar contains the following menu items: DASHBOARD, SYSTEM SETTINGS (with sub-items: General Settings, Management Port, Power Settings, Notifications, and Maintenance), HOST CONNECTIVITY (with sub-items: Overview and iSCSI Ports), STORAGE MANAGEMENT (with sub-items: Disks, Pools, Volumes, LUN Mappings, and SSD Cache), and DATA PROTECTION (with sub-items: Snapshots and Remote Replications). The 'Maintenance' item is highlighted. The main content area shows the 'System Information' page, which includes tabs for 'System Information', 'Update', 'Firmware Synchronization', 'System Identification', and 'Reset'. The 'System Information' tab is active, displaying the following information:

| Item | Information |
|-------------------------------|----------------------------|
| Firmware Version | 1.2.0 (build 201707171500) |
| SAS IOC Firmware Version | 07.00.01.00 |
| SAS Expander Firmware Version | 1000 |

Controller 2

| Item | Information |
|-------------------------------|----------------------------------|
| Controller Serial Number | |
| CPU | Intel(R) Processor D1500 2 Cores |
| Memory | 8 GB |
| Host Card Slot 1 | Empty |
| Host Card Slot 2 | Empty |
| Firmware Version | 1.2.0 (build 201707171500) |
| SAS IOC Firmware Version | 07.00.01.00 |
| SAS Expander Firmware Version | 1000 |

Backplane

| Item | Information |
|-------------------------|-------------|
| Backplane Serial Number | |
| Backplane ID | QW424 |
| MCU Version | 1.2.0 |

4.2. Documentation Feedback

QSAN is committed to providing documentation that meets and exceeds your expectations. To help us improve the documentation, email any errors, suggestions, or comments to docsfeedback@qsan.com.

When submitting your feedback, include the document title, part number, revision, and publication date located on the front cover of the document.

Appendix

End-User License Agreement (EULA)

Please read this document carefully before you use our product or open the package containing our product.

YOU AGREE TO ACCEPT TERMS OF THIS EULA BY USING OUR PRODUCT, OPENING THE PACKAGE CONTAINING OUR PRODUCT OR INSTALLING THE SOFTWARE INTO OUR PRODUCT. IF YOU DO NOT AGREE TO TERMS OF THIS EULA, YOU MAY RETURN THE PRODUCT TO THE RESELLER WHERE YOU PURCHASED IT FOR A REFUND IN ACCORDANCE WITH THE RESELLER'S APPLICABLE RETURN POLICY.

General

QSAN Technology, Inc. ("QSAN") is willing to grant you ("User") a license of software, firmware and/or other product sold, manufactured or offered by QSAN ("the Product") pursuant to this EULA.

License Grant

QSAN grants to User a personal, non-exclusive, non-transferable, non-distributable, non-assignable, non-sub-licensable license to install and use the Product pursuant to the terms of this EULA. Any right beyond this EULA will not be granted.

Intellectual Property Right

Intellectual property rights relative to the Product are the property of QSAN or its licensor(s). User will not acquire any intellectual property by this EULA.

License Limitations

User may not, and may not authorize or permit any third party to: (a) use the Product for any purpose other than in connection with the Product or in a manner inconsistent with the design or documentations of the Product; (b) license, distribute, lease, rent, lend, transfer, assign or otherwise dispose of the Product or use the Product in any commercial hosted or service bureau environment; (c) reverse engineer, decompile, disassemble or attempt to discover the source code for or any trade secrets related to the Product, except and only to the extent that such activity is expressly permitted by applicable law notwithstanding this limitation; (d) adapt, modify, alter, translate or create any derivative works of the Licensed Software; (e) remove, alter or obscure any copyright notice or other proprietary rights notice on the Product; or (f) circumvent or attempt to circumvent any methods employed by QSAN to control access to the components, features or functions of the Product.

Disclaimer

QSAN DISCLAIMS ALL WARRANTIES OF PRODUCT, INCLUDING BUT NOT LIMITED TO ANY MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, WORKMANLIKE EFFORT, TITLE, AND NON-INFRINGEMENT. ALL PRODUCTS ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND. QSAN MAKES NO WARRANTY THAT THE PRODUCT WILL BE FREE OF BUGS, ERRORS, VIRUSES OR OTHER DEFECTS.

IN NO EVENT WILL QSAN BE LIABLE FOR THE COST OF COVER OR FOR ANY DIRECT, INDIRECT, SPECIAL, PUNITIVE, INCIDENTAL, CONSEQUENTIAL OR SIMILAR DAMAGES OR LIABILITIES WHATSOEVER (INCLUDING, BUT NOT LIMITED TO LOSS OF DATA, INFORMATION, REVENUE, PROFIT OR BUSINESS) ARISING OUT OF OR RELATING TO THE USE OR INABILITY TO USE THE PRODUCT OR OTHERWISE UNDER OR IN CONNECTION WITH THIS EULA OR THE PRODUCT, WHETHER BASED ON CONTRACT, TORT (INCLUDING NEGLIGENCE), STRICT LIABILITY OR OTHER THEORY EVEN IF QSAN HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Limitation of Liability

IN ANY CASE, QSAN'S LIABILITY ARISING OUT OF OR IN CONNECTION WITH THIS EULA OR THE PRODUCT WILL BE LIMITED TO THE TOTAL AMOUNT ACTUALLY AND ORIGINALLY PAID BY CUSTOMER FOR THE PRODUCT. The foregoing Disclaimer and Limitation of Liability will apply to the maximum extent permitted by applicable law. Some jurisdictions do not allow the exclusion or limitation of incidental or consequential damages, so the exclusions and limitations set forth above may not apply.

Termination

If User breaches any of its obligations under this EULA, QSAN may terminate this EULA and take remedies available to QSAN immediately.

Miscellaneous

- QSAN reserves the right to modify this EULA.
- QSAN reserves the right to renew the software or firmware anytime.
- QSAN may assign its rights and obligations under this EULA to any third party without condition.
- This EULA will be binding upon and will inure to User's successors and permitted assigns.
- This EULA shall be governed by and constructed according to the laws of R.O.C. Any disputes arising from or in connection with this EULA, User agree to submit to the jurisdiction of Taiwan Shilin district court as first instance trial.