

Storage Performance Tuning Guide

Best Practice

July 2022

ANNOUNCEMENT

Copyright

© Copyright 2022 QSAN Technology, Inc. All rights reserved. No part of this document may be reproduced or transmitted without written permission from QSAN Technology, Inc.

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, QSAN.com, XCubeFAS, XCubeSAN, XCubeNXT, XCubeNAS, XCubeDAS, XEVO, SANOS, and QSM are trademarks or registered trademarks of QSAN Technology, Inc.
- Microsoft, Windows, Windows Server, and Hyper-V are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries.
- Linux is a trademark of Linus Torvalds in the United States and/or other countries.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Mac and OS X are trademarks of Apple Inc., registered in the U.S. and other countries.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.
- VMware, ESXi, and vSphere are registered trademarks or trademarks of VMware, Inc. in the United States and/or other countries.
- Citrix and Xen are registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or 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.

TABLE OF CONTENTS

Announcement	i
Notices	iv
Preface	v
Executive Summary	v
Audience	v
Technical Support	v
Information, Tip, and Caution	vi
1. Performance Tuning Guide	1
1.1. Prerequisites	1
1.2. Guidelines for Hosts.....	2
1.3. Guidelines for Configuring Storage Pools.....	3
1.4. Guidelines for Configuring iSCSI Connections	4
1.5. Guidelines for Configure FC Connections.....	15
2. Test Results and Use Cases	20
2.1. Performance Results.....	20
2.2. Case1 – High IOPS with Low Latency.....	22
2.3. Case2 – Useful Onboard Ports.....	22
2.4. Case3 – Extreme High Throughput.....	22
3. Conclusion	24
4. Appendix	25
4.1. Apply To	25
4.2. Reference.....	25

FIGURES

Figure 1-1	Storage Connection Topology	1
Figure 1-2	iSCSI Ports in XEVO	5
Figure 1-3	iSCSI Ports in SANOS	5
Figure 1-4	Network Interfaces in QSM	6
Figure 1-5	Adjust the Read Ahead Buffer of the RAID Volume	10
Figure 1-6	Adjust the Receive Buffer of TCP/IP	10
Figure 1-7	Assign the PCIE Device from the VMware ESXi Server Step 1.....	12
Figure 1-8	Assign the PCIE Device from the VMware ESXi Server Step 2.....	12
Figure 1-9	Assign the PCIE Device from the VMware ESXi Server Step 3.....	13
Figure 1-10	Assign the PCIE Device from the VMware ESXi Server Step 4	13
Figure 1-11	Assign the PCIE Device from the VMware ESXi Server Step 5	14
Figure 1-12	FC Ports in XEVO	16
Figure 1-13	FC Ports in SANOS.....	16
Figure 1-14	Adjust the Read Ahead Buffer of the RAID Volume	19
Figure 2-1	Performance of Random Read 4K	21
Figure 2-2	Performance of Random Write 4K	21

NOTICES

Information contained in this document 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

Executive Summary

This document provides technical guidance for setting up iSCSI and FC (Fibre Channel) connections with QSAN XCubeFAS, XCubeSAN, XCubeNXT, and XCubeNAS series products under Windows, Unix-Like OS and VMware environments.

We will show the performance of the recommended configuration and highlight the use cases of different applications at the end of this document. It will help users choose QSAN storage to solve their challenges. As a multi-purpose storage, no matter what the hardware or software parameters in the environment, it is recommended to use the appropriate configuration to meet the corresponding requirements.

Audience

This document is applicable for QSAN customers and partners who are interested in learning iSCSI and FC implementation. It assumes the reader is familiar with QSAN products and has general IT experience, including knowledge as a system or network administrator. If there is any question, please refer to the user manuals of products, or contact QSAN support for further assistance.

Technical Support

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

- Via the Web: https://www.qsan.com/technical_support
- Via Telephone: +886-2-77206355
- (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, Summer time: 09:30 - 01:00)

- Via Email: support@qsan.com

Information, Tip, and Caution

This document 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. PERFORMANCE TUNING GUIDE

This chapter provides technical guidance for setting up iSCSI and FC (Fibre Channel) connections to achieve better performance in Windows, Unix-Like OS, and VMware environments.

1.1. Prerequisites

The following figure describes a storage connection topology. Taking XF2026 or XS5226 / XS3226 / XS1226 as a storage example, the figure includes two hosts, management, iSCSI, FC connections, and one computer for management.

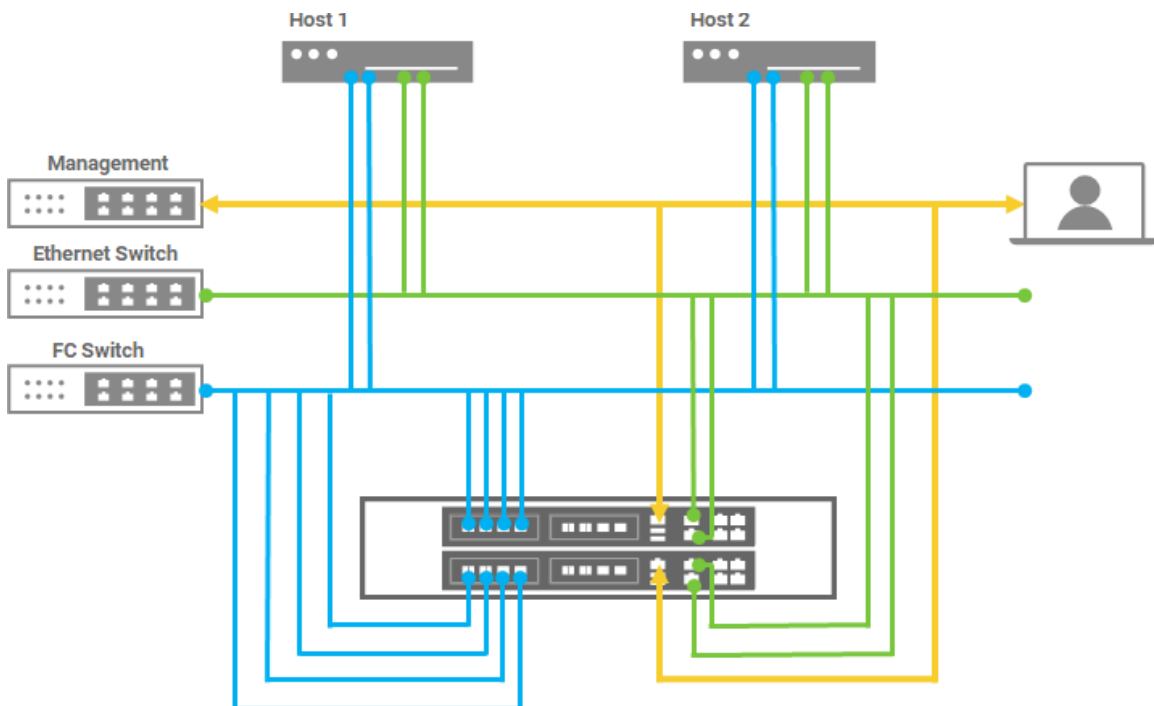


Figure 1-1 Storage Connection Topology



TIP

Ensure that all network card ports used for iSCSI connections on the hosts can ping and be pinged to the storage iSCSI data ports.

Several guidelines for storage connection topology are as follows.

1. For XCubeFAS and XCubeSAN series products, it is recommended to use different IP segments for **iSCSI** connectivity and **Management** connections. The Management port (green line, the port above the USB ports) can only be used for management, and data I/O cannot be passed through this port.
2. For XCubeNXT and XCubeNAS series products, there is no such restriction, because each port can be the management or data I/O ports, or they can be used in combination.
3. Please also consider configuring different network segments for each network card port on the host and iSCSI data port on the storage for later differentiation and troubleshooting.
4. The host card on each controller is optional. Add FC or iSCSI host cards in one or two slots as needed.



INFORMATION

Due to XF2026, XS5200 / XS3200 / XS1200, XN8000D, XN8024R / XN8016R / XN8012S / XN7024R / XN7016R hardware design limitations, Slot 2 has lower bandwidth and slower throughput. XF3126 is a new architecture without this limitation.

-
5. The onboard 10GbE ports (blue line the two RJ45 ports on the right-hand side) can be used for iSCSI connection or remote replication for DR (Disaster Recovery) purpose.

1.2. Guidelines for Hosts

This section introduces the host guidelines. Nowadays, the mainstream operating systems include drivers for most popular HBAs (Host Bus Adapters) by default, but it is still recommended to update the HBA drivers installed on the host to the latest version to prevent incompatibility.



TIP

Update the driver version of the HBA (Host Bus Adapter) on the host side, or any firmware / driver of the relevant hardware involved in the environment, such as Ethernet or Fibre Channel switches. These can prevent any possible abnormalities caused by the old firmware / software version.

1.2.1. High-Performance Local Drive on the Host

If you are going to test performance through copy and paste operations, please make sure that the local drive installed on the host possesses high performance, because those operations are also related to the read and write performance of the local drive.

For example, if the local drive only has a read performance of 100MB/s, while a RAID volume can provide write throughput of up to 1,000MB/s. When copying data from the local drive to the storage RAID volume, the test result will be around 100MB/s due to the bottleneck of the local drive.

1.3. Guidelines for Configuring Storage Pools

This section introduces the storage configuration guidelines. QSAN will provide the best configuration by default when creating a RAID pool to try to get the best performance. If there is no special requirement, please do not modify any of the following parameters.

1. Using Thick Provisioning to create a pool, its performance is better than Thin. If you have to use Thin, it is OK. Please make sure that you have carefully read the policy of Thin Provisioning, otherwise data inconsistency may occur.
2. When creating a pool, make sure to enable the Disk Write Cache, Disk Read-ahead, and Disk Command Queuing for better performance. These settings are enabled by default.
3. When creating a volume, ensure to enable the Write-back Cache Mode for better performance.
4. [Optional] Enable the Video Editing Mode when your application is in video editing environment. This option provides a more stable performance figure without high and low peaks but slower in average.

5. [Optional] Disable the Fast RAID Rebuild for better performance. The setting is disabled by default.
6. [Optional] DO NOT take snapshots when verifying performance. Copy-on-write when taking a snapshot adds an additional burden and impairs write performance.
7. [Optional] There is an important concept. If the overall performance of the disk drives (theoretically) exceeds the front-end port threshold, try to divide them into multiple pools and evenly distribute them to the two controllers. For example, the sequential read performance exceeds 6,400MB/s if a 4-port 16Gb FC host card is used. Or it exceeds 4,000MB/s if a 4-port 10GbE iSCSI host card is used.
8. [Optional] If you do not need to use LUN Masking, keep the default setting "*", which means the volume can be accessed by any host.

**TIP**

If you plan to use this volume in a VMware ESXi environment, please **DO NOT** set it to 4K block size. So far (February 2020) VMware does not support 4K block size external storage.

**INFORMATION**

For XCubeNXT and XCubeNAS users, if the above settings do not appear in QSM, they can be ignored.

1.4. Guidelines for Configuring iSCSI Connections

This section introduces iSCSI connection guidelines, including storage configuration, network switch, and host configuration. Although the setting method is different from XEVO, SANOS, and QSM, the following screenshots describe the configuration separately.

1.4.1. Configure iSCSI Settings in Storage

This section introduces guidelines for configuring iSCSI settings in storage.

1. Enter the iSCSI port page to configure iSCSI network settings. In XEVO, select the **System** tab and the **Data Ports** subtab to configure iSCSI data ports. In SANOS, select the **iSCSI Ports** function submenu to show the information of iSCSI ports. In QSM, select the **Network** function submenu in the Control Panel to display the network interfaces.

CTRL	Interface	Location	Port	Status	LAG	IP Address	Gateway
1	iSCSI (10GB)	Onboard	LAN1	1 Gb/s	N/A	10.10.1.1	
1	iSCSI (10GB)	Onboard	LAN2	1 Gb/s	N/A	10.10.1.2	
2	iSCSI (10GB)	Onboard	LAN1	1 Gb/s	N/A	10.10.1.3	
2	iSCSI (10GB)	Onboard	LAN2	1 Gb/s	N/A	10.10.1.4	

Figure 1-2 iSCSI Ports in XEVO

Controller 1									
	Location	Port Name	Status	LAG	VLAN ID	IP Address	Gateway	Jumbo Frame	MAC Address
▼	Onboard	LAN1 (10Gb)	Down	N/A	N/A	10.10.1.101		Disabled	00:13:78:d4:23:42
▼	Onboard	LAN2 (10Gb)	1 Gb/s	N/A	N/A	10.10.1.102		Disabled	00:13:78:d4:23:43
Controller 2									
	Location	Port Name	Status	LAG	VLAN ID	IP Address	Gateway	Jumbo Frame	MAC Address
▼	Onboard	LAN1 (10Gb)	Down	N/A	N/A	10.10.1.103		Disabled	00:13:78:d4:23:4c
▼	Onboard	LAN2 (10Gb)	1 Gb/s	N/A	N/A	10.10.1.104		Disabled	00:13:78:d4:23:4d

Figure 1-3 iSCSI Ports in SANOS

Cluster IP

Edit Delete

No.	Controller 1	Controller 2	IP address	Mask
1	LAN1	LAN1	192.168.170.23	255.255.0.0
2	LAN2	LAN2	-	-
3	LAN3	LAN3	-	-

Interface

Controller 1

Edit

Link	Interfaces	Jumbo frame	VLAN	IP address	Gateway	Speed	MAC address
●	LAN1 (1 GbE)	1500	0	192.168.170.11	192.168.128.254	1 Gbps	00:13:78:DC:2...
●	LAN2 (10 G...	1500	0	10.10.1.201	10.10.1.254	Down	00:13:78:DC:2...
●	LAN3 (10 G...	1500	0	10.10.1.202	10.10.1.254	Down	00:13:78:DC:2...

Controller 2

Edit

Link	Interfaces	Jumbo frame	VLAN	IP address	Gateway	Speed	MAC address
●	LAN1 (1 GbE)	1500	0	192.168.170.22	192.168.128.254	1 Gbps	00:13:78:DC:2...
●	LAN2 (10 G...	1500	0	10.10.1.203	10.10.1.254	Down	00:13:78:DC:2...
●	LAN3 (10 G...	1500	0	10.10.1.204	10.10.1.254	Down	00:13:78:DC:2...

Figure 1-4 Network Interfaces in QSM

**TIP**

First check your environment and decide how to set it up. Use a 10GbE link if necessary. The onboard 10GbE iSCSI (RJ45) ports are compatible with 1GbE links, and all ports can be connected to the same switch.

2. Ensure that each network port is used for I/O transmission and operating at the expected link speed. You can view it from the performance monitoring page of the Web UI.
3. Make sure you have enough hosts to request the same high performance as the controller port. For example, if there are 4 10GbE ports in the controller, at least 4 10GbE ports must be prepared from the host. Expect a 10GbE port can reach a maximum performance about 1,000MB/s
4. [Optional] Configure **Jumbo Frame** according to the settings in your environment. Enabling Jumbo Frame will reduce the CPU and network usage. Please configure the appropriate size. If you enable Jumbo Frame from the controller, make sure that all ports (including switches) are enabled with the same size as the controller port (also called **MTU** (Maximum Transmission Unit)).
5. [Optional] Configure **VLAN** (Virtual Local Area Network) according to the settings in your environment. By dividing different VLANs, devices in VLANs can communicate directly, but

they cannot communicate directly between VLANs. The advantage of VLAN is to limit the broadcast domain, enhance security, improve the robustness of the network, and flexibly construct virtual workgroups. Please configure the appropriate number. If you set VLAN ID from the controller, make sure that all ports (including switches) are set with the same number as the controller port.

6. [Optional] You can also consider changing the iSCSI **Entity Name**. After applying the setting, all iSCSI target names will be changed to the defined names. Of course, to make it easier to understand.
7. It is **NOT** recommended to use **Trunking** or **LACP** (Link Aggregation Control Protocol) unless you plan to connect this storage to a multi-client topology, such as more than 10 clients or hosts. Trunking or LACP will help simplify connections, but enabling MPIO (Multipath I/O) connections on the host side is sufficient. If Trunking or LACP is enabled on the controller, please ensure that the switch is configured with the same settings to meet the requirements. All connected ports on the switch need to be in the same network group.

1.4.2. Configure Network Settings in Ethernet Switch

This section introduces guidelines for Ethernet switches. Although the settings of each switch manufacturer may be different, the concept is the same. We usually check the following options.

1. **Jumbo Frame:** Set the same number on the NIC (Network Interface Card) port according to the configuration in the environment.
2. **Flow Control:** Sometimes must be set to ON, but we did encounter some situations where it must be set to OFF, please check the environment and verify the performance after the changes.
3. **Trunking / (Link Aggregation Control Protocol):** You must first be aware of the existing network configuration. If you use LACP, please use the same protocol to apply. These settings are suitable for multi-client topologies.



TIP

If you still encounter slow performance after following the above suggestions, please enable port mirroring on the Ethernet switch. Connect the network cable from the Windows client to the mirror port of the switch. Install the [Wireshark](#) software and capture network packets in the performance test. And then share the result files to QSAN technical support team. We will help clarify the reason for the slowness.

1.4.3. Configure iSCSI Settings in Windows

This section presents guidelines for configuring iSCSI settings in Windows. For better performance, we usually check the following options.

1. Make sure that each network port is used for I/O transmission. You can check the performance tab it in the **Task Manager** of Windows Server.
2. When logging in to the iSCSI target from the Windows Server iSCSI Initiator, you can specify the source IP address to ensure that each NIC port on the host side will be used.

After confirming the above configuration, you may experience poor performance when testing performance via iSCSI connection. If encountered, please try to adjust the following parameters in the HBA driver.

1. Access the **Device Manager** in the **Network** adapters, right-click the port for performance testing
2. Select the **Properties**, go to the **Advanced**, find the **RSS Queue** (different brands may have different terms), and adjust the value from default 8 to 2.
3. Find the **Receive Buffers** and adjust the value to the maximum.
4. Find the **Transmit Buffers** and adjust the value to the maximum.
5. Find the **Interrupt Moderation Rate** and adjust the value to OFF.
6. If the above adjustments did not help, please further open a Command Line and type the following command.

```
C:\> netsh int tcp set global autotuninglevel=restricted  
Or  
C:\> netsh int tcp set global autotuninglevel=highlyrestricted
```

7. After setting the above parameters, you need to log in to the connected iSCSI session again. If it does not work, you may also need to restart the host.

**TIP**

Please verify the performance immediately after changing any of the above adjustments to find one that suits the user's environment, without adjusting one by one.

1.4.4. Configure iSCSI Settings in Unix-Like OS

This section presents guidelines for configuring iSCSI settings in Unix-Like OS. Unix-Like OS may not support specifying the source IP address when logging in to an iSCSI target, so you may need to configure a different network segment for each NIC port on the host.

**TIP**

Unix-like operating systems cannot specify the source NIC port to log in to the defined target iSCSI portal. It is possible that all IP addresses are on the same network segment. The operating system always uses the same NIC port as the source to log in to a different iSCSI portal. Therefore, the overall performance of the network card will be limited.

If you experience poor performance when testing performance via iSCSI connection, please try to adjust the following parameters on the host.

1. Adjust the RA (Read Ahead) buffer of the RAID volume through the **blockdev** command. Set a higher value for each RAID volume, such as 4,096 or 8,192. If the RAID volume is a multipath device, please adjust the dm-x device.

```
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
```

Figure 1-5 Adjust the Read Ahead Buffer of the RAID Volume

2. Adjust the **Receive Buffer** of TCP/IP. Set it to 524,284 or higher (twice the value).

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

Figure 1-6 Adjust the Receive Buffer of TCP/IP

3. Disable HT (Hyper Threading) of the CPU Configuration in the BIOS setting.
4. After setting the above parameters, you need to log in to the connected iSCSI session again. If it does not work, you may also need to restart the host.



TIP

Please verify the performance immediately after changing any of the above adjustments to find one that suits the user's environment, without adjusting one by one.

1.4.5. Configure iSCSI Settings in VMware

This section presents guidelines for configuring iSCSI settings in VMware. The same as Unix-Like OS, VMware may not support specifying the source IP address when logging in to an iSCSI target, so you may need to configure a different network segment for each NIC port on the host.



TIP

VMware ESXi cannot specify the source NIC port to log in to the defined target iSCSI portal. It is possible that all IP addresses are on the same network segment. The operating system always uses the same NIC port as the source to log in to a different iSCSI portal. Therefore, the overall performance of the network card will be limited.

If you encounter any problems while using VMware, please check the following options.

1. If you plan to use this volume in a VMware ESXi environment, please **DO NOT** set it to 4K block size. So far (February 2020) VMware does not support 4K block size external storage.
2. If MPIO is enabled in the VMware ESXi server, it is recommended to set it to **Round Robin**. Please remember to adjust the IOPS value from default 1000 to 1. (from VMware Knowledge Base [#2069356](#))
3. If you are going to add an additional HBA, make sure to configure it correctly.
 - A. Assign the PCIE device from the VMware ESXi server to the VM (Virtual Machine).
 - B. After this operation, the ESXi server needs to be rebooted.

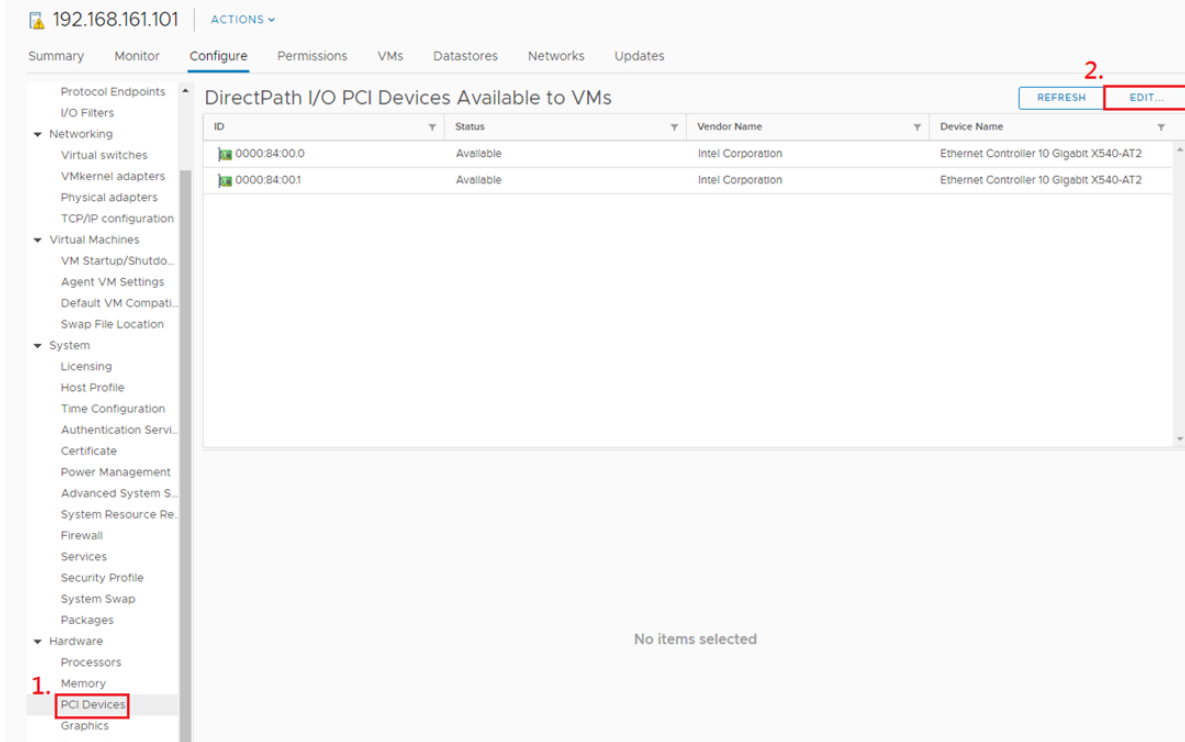


Figure 1-7 Assign the PCIe Device from the VMware ESXi Server Step 1

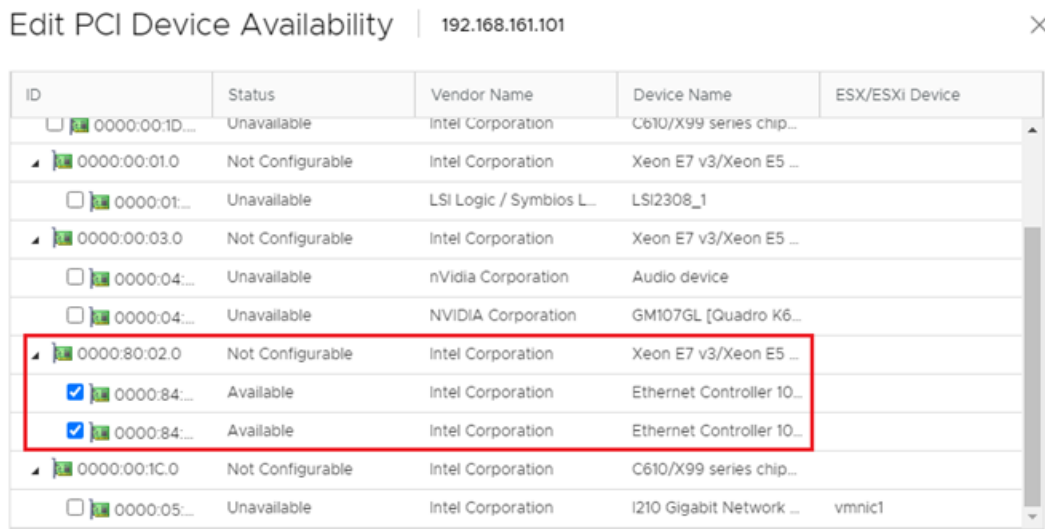


Figure 1-8 Assign the PCIe Device from the VMware ESXi Server Step 2

C. Access the virtual machine and edit the device.

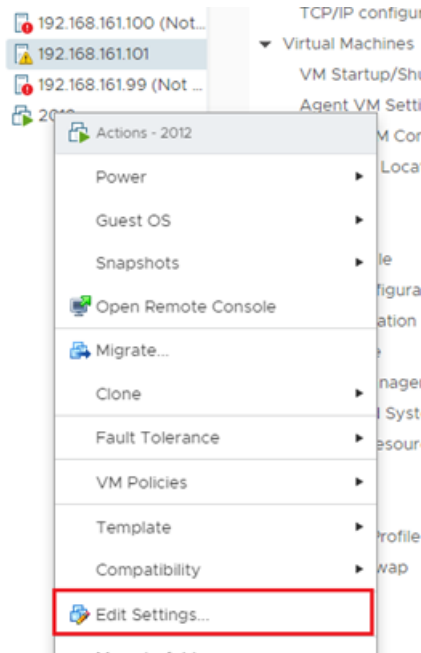


Figure 1-9 Assign the PCIe Device from the VMware ESXi Server Step 3

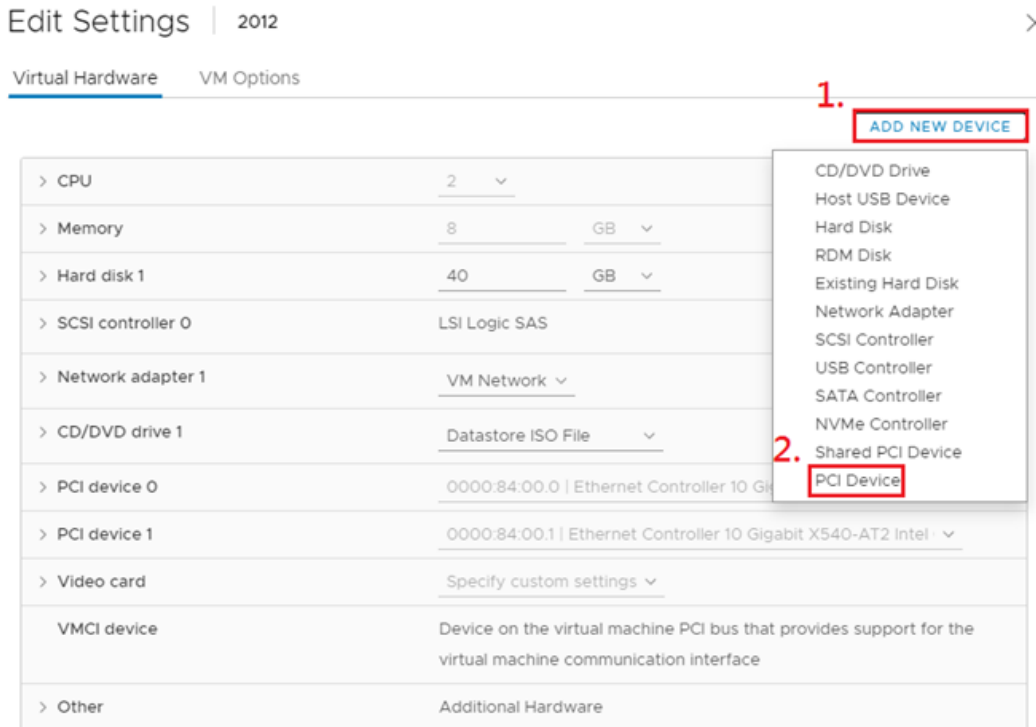


Figure 1-10 Assign the PCIe Device from the VMware ESXi Server Step 4

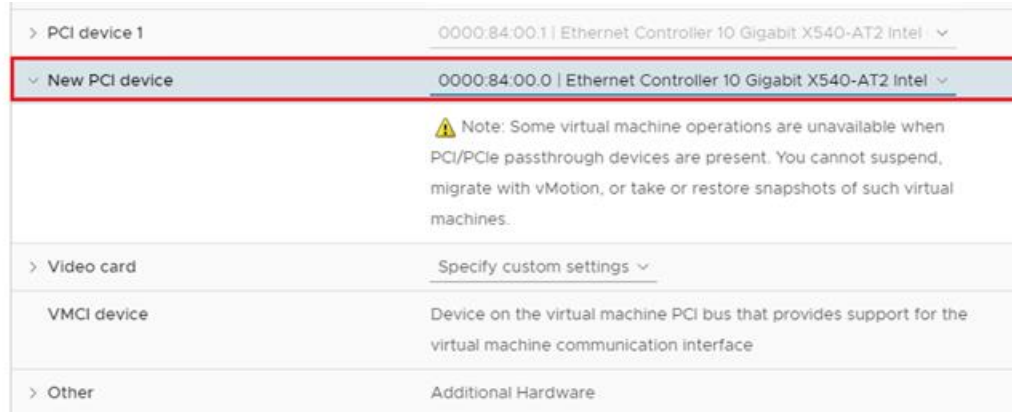


Figure 1-11 Assign the PCIe Device from the VMware ESXi Server Step 5

D. Done.

- If you encounter a problem regarding the disconnection between VMware ESXi server and storage, VMware support may ask for the function of **Delayed ACK**.



TIP

QSAN storage supports **Delayed ACK**, but we are the receiver of this feature. QSAN has our own mechanism to control (queue buffer) and then confirm with the initiator (client) to optimize performance. Disabling ACK on the ESXi server basically does not affect the iSCSI service.

- In VMware Knowledge Base [#2113956](#), it mentioned that an ESXi host loses connectivity to a VMFS datastore with VAAI ATS. The recommended resolution from VMware is to disable VAAI ATS heartbeat on ALL hosts sharing the datastore where these errors are seen.

QSAN had conducted thorough testing in the lab and concluded that such issue would lead to data inconsistency which may corrupt data before users are aware of this compatibility issue. The proactive approach from QSAN is to provide solution from storage side so customer could upgrade the firmware in storage side before upgrading to VMware 7.0 U2. It will prevent this issue while still maintaining existing environment and all VMware advanced features without interruption.



TIP

This solution has been added to the latest firmware in QSAN storage. Please refer to the [Apply To](#) section for more details.

1.5. Guidelines for Configure FC Connections

This section introduces Fibre Channel connection guidelines, including storage configuration, Fibre Channel switch, and host configuration. Although the setting method is different from XEVO, SANOS, and QSM, the following screenshots describe the configuration separately.

1.5.1. Configure FC Settings in Storage

This section introduces guidelines for configuring iSCSI settings in storage.

1. Enter the Fibre Channel port page to configure iSCSI network settings. In XEVO, select the **System** tab and the **Data Ports** subtab to configure FC data ports. In SANOS, select the **Fibre Channel Ports** function submenu to show the information of FC ports.

Fibre Channel Ports

CTRL	Interface	Location	Port	Status	Topology	WWNN/WWPN
1	FC (16GB)	Slot1	FC1	16GB/s	Point-to-Point	WWNN: 2000001378D61D5 WWPN: 2100001378D61D5
1	FC (16GB)	Slot1	FC2	16GB/s	Point-to-Point	WWNN: 2000001378D61D5 WWPN: 2200001378D61D5
1	FC (16GB)	Slot1	FC3	Down	NA	WWNN: 2000001378D61D5 WWPN: 2300001378D61D5
1	FC (16GB)	Slot1	FC4	Down	NA	WWNN: 2000001378D61D5 WWPN: 2400001378D61D5
2	FC (16GB)	Slot1	FC1	16GB/s	Point-to-Point	WWNN: 2000001378D61D5 WWPN: 2900001378D61D5
2	FC (16GB)	Slot1	FC2	16GB/s	Point-to-Point	WWNN: 2000001378D61D5 WWPN: 2A00001378D61D5
2	FC (16GB)	Slot1	FC3	Down	NA	WWNN: 2000001378D61D5 WWPN: 2B00001378D61D5
2	FC (16GB)	Slot1	FC4	Down	NA	WWNN: 2000001378D61D5 WWPN: 2C00001378D61D5

Figure 1-12 FC Ports in XEVO

Controller 1

	Location	Name	Status	Change Topology	WWNN/WWPN	Loss of Signal	Loss of Sync	Link Failure	Invalid CRC
▼	Slot1	FC1 (16Gb)	16Gb/s	Point-to-Point	WWNN: 2000001378124770 WWPN: 2100001378124770	0	655	0	0
▼	Slot1	FC2 (16Gb)	Down	NA	WWNN: 2000001378124770 WWPN: 2200001378124770	0	0	0	0
▼	Slot1	FC3 (16Gb)	Down	NA	WWNN: 2000001378124770 WWPN: 2300001378124770	0	0	0	0
▼	Slot1	FC4 (16Gb)	Down	NA	WWNN: 2000001378124770 WWPN: 2400001378124770	0	0	0	0

Controller 2

	Location	Name	Status	Change Topology	WWNN/WWPN	Loss of Signal	Loss of Sync	Link Failure	Invalid CRC
▼	Slot1	FC1 (16Gb)	16Gb/s	Point-to-Point	WWNN: 2000001378124770 WWPN: 2900001378124770	0	561	0	0
▼	Slot1	FC2 (16Gb)	Down	NA	WWNN: 2000001378124770 WWPN: 2A00001378124770	0	0	0	0
▼	Slot1	FC3 (16Gb)	Down	NA	WWNN: 2000001378124770 WWPN: 2B00001378124770	0	0	0	0
▼	Slot1	FC4 (16Gb)	Down	NA	WWNN: 2000001378124770 WWPN: 2C00001378124770	0	0	0	0

Clear All Counters

Figure 1-13 FC Ports in SANOS

2. [Optional] Configure **Topology** according to the existing FC environment. The default setting is P2P (Point-to-Point). If the peer (FC switch or FC HBA) uses the automatic connection mode, it will be adjusted to the same automatically.



INFORMATION

16Gb FC only supports **Point-to-Point** topology. If using **Loop** mode, please adjust the connection speed to 8G or 4G. These changes require a restart of the storage.

-
3. [Optional] Configure **Link Speed** according to the existing FC environment. The default is Automatic. If the peer end uses Automatic as well, the highest speed will be negotiated.

1.5.2. Configure Network Settings in FC Switch

This section introduces guidelines for FC switches. Although the settings of each switch manufacturer may be different, the concept is the same. We usually check the following options.

1. Set the FC **Topology** and **Link Speed** to Automatic settings to establish the connection correctly.
2. [Optional] If necessary, configure the **ZONE** settings to ensure that the relevant FC ports are included in the same ZONE group so that the host can detect the FC LUN.



TIP

If the FC LUN cannot be detected or the host is abnormal, please abandon the ZONE configuration. Starting from the default settings, and then add the preferred ZONE configuration one by one.

1.5.3. Configure FC Settings in Windows

This section presents guidelines for configuring FC settings in Windows. For better performance, we usually check the following options.

1. Ensure that the driver of the installed FC HBA has been updated to the latest version, and ensure that the FC connection has been connected at the expected link speed.
2. [Optional] If using Marvell QLogic FC HBA, execute "regedit" under Windows Server to open the registry editor and access "HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ql2300i\Parameters\Device", adjust the value in **DriverParameter** from the default value to "qd=255", and then restart the host.

1.5.4. Configure FC Settings in Unix-Like OS

This section presents guidelines for configuring FC settings in Unix-Like OS.

1. Ensure that the driver of the installed FC HBA has been updated to the latest version, and ensure that the FC connection has been connected at the expected link speed.
2. [Optional] If MPIO is enabled on the host, it is recommended to set `rr_min_io` value to "1" in the `multipath.conf` file.

If you experience poor performance when testing performance via FC connection, please try to adjust the following parameters on the host.

1. Adjust the RA (Read Ahead) buffer of the RAID volume through the **blockdev** command. Set a higher value for each RAID volume, such as 4,096 or 8,192. If the RAID volume is a multipath device, please adjust the `dm-x` device.

```

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

```

Figure 1-14 Adjust the Read Ahead Buffer of the RAID Volume

2. Disable HT (Hyper Threading) of the CPU Configuration in the BIOS setting.
3. After setting the above parameters, you need to reset the connected FC session again. If it does not work, you may also need to restart the host.



TIP

Please verify the performance immediately after changing any of the above adjustments to find one that suits the user's environment, without adjusting one by one.

1.5.5. Configure FC Settings in VMware

This section presents guidelines for configuring FC settings in VMware. If you encounter any problems while using VMware, please check the following options.

1. If you plan to use this volume in a VMware ESXi environment, please **DO NOT** set it to 4K block size. So far (February 2020) VMware does not support 4K block size external storage.
2. If MPIO is enabled in the VMware ESXi server, it is recommended to set it to **Round Robin**. Please remember to adjust the IOPS value from default 1000 to 1. (from VMware Knowledge Base [#2069356](#))

2. TEST RESULTS AND USE CASES

This chapter provides test results and use cases for verification.

2.1. Performance Results

We verified the maximum IOPS through the Vdbench benchmark tool with a latency of less than 1ms. This is the environment and results.

Test Equipments and Configurations

- Storage
 - Model: XCubeFAS XF2026D
 - Memory: 32GB (4 x 8GB) per controller
 - Firmware: 1.0.0
 - SAS SSD: 24 x Seagate Nytro 3530, XS3200LE10003, 3.2TB, SAS 12Gb/s
 - Pools: 2 x (13 x SSDs per Pool for RAID 5, RAID 6)
2 x (12 x SSDs per Pool for RAID 10)
 - Volumes: 4 x 100GB in Pool 1 (Ctrl 1); 4 x 100GB in Pool 2 (Ctrl 2)
 - Volume Block Size: 4,096 Bytes
- Server
 - Model: 1 x Dell E25S; 1 x HP Z840
 - 16Gb FC HBA: Marvell QLogic QLE2694 / QLE2672
 - OS: Windows Server 2012 R2
- Vdbench
 - Version: 5.04.06
 - Workers: 2 x 4 (1 Worker to 1 Volume)
 - Outstanding I/Os: 128
 - Xfersize: 4K 8K 32K 64K
 - I/O rates: 10 - 120
 - Reporting Interval: 1 sec
 - Warmup period: 5 sec
 - Elapsed Time: 30 sec per I/O rate

Test Results

The following are the performance results of random read and write in RAID 5, RAID 6, and RAID 10.

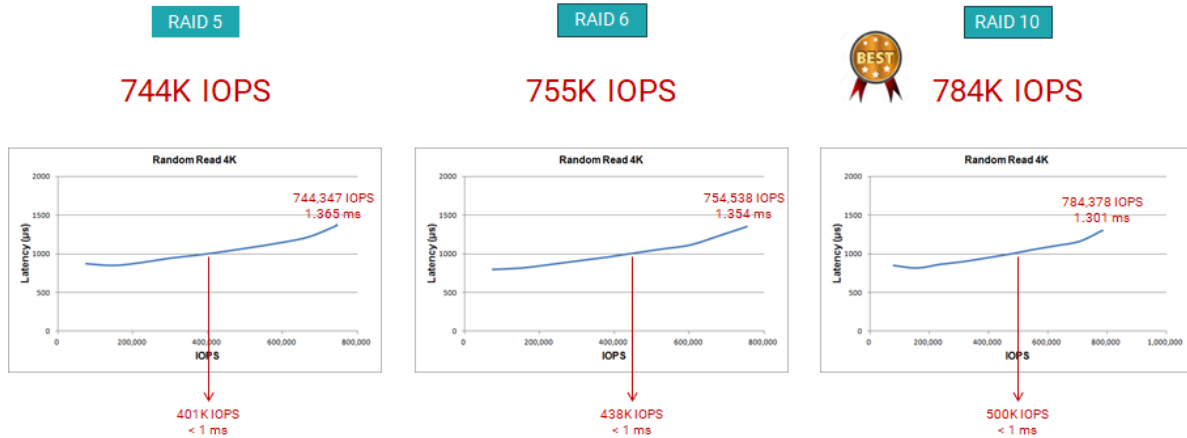


Figure 2-1 Performance of Random Read 4K

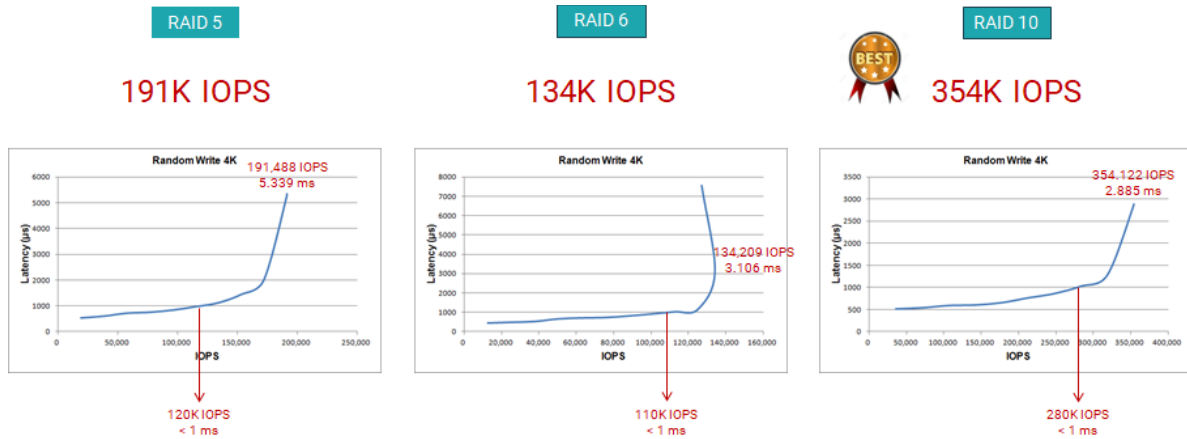


Figure 2-2 Performance of Random Write 4K

It can be seen from the above results that under RAID 10 level, it provides good performance of about 500K IOPS in random read 4K under 1ms, and the maximum is about 784K IOPS. Similarly, it provides about 280K IOPS in random write 4K under 1ms, and the maximum is about 384K IOPS.

2.2. Case1 – High IOPS with Low Latency

In virtualization or VDI (Virtual Desktop Infrastructure) applications, high IOPS and low latency are required to run smoothly. XCubeFAS can provide very good performance as shown in the previous section. Some guidelines are as follows.

- Use **Fibre Channel** if such requirements are required.
- Configure the RAID level to **RAID 10** for small and random IOPS environments.
- Make sure to follow the relevant configuration and parameters recommended in the previous chapter.

This high IOPS with low latency is suitable for the following situations.

- VDI or any virtualized environments
- SQL database
- Media editing
- IoT (Internet of Things)

It is recommended to a higher-level storage array, such as AFA (All Flash Array) series XF3126 or XF2026.

2.3. Case2 – Useful Onboard Ports

XCubeSAN series products come with onboard 2-port 10GbE iSCSI (RJ45). These ports are only used for the iSCSI protocol for further access. In addition, if DR (Disaster Recovery) solution is required, these are also very useful ports when performing the remote replication between the two XCubeSAN units. There is similar concept for XCubeFAS, XCubeNXT, and XCubeNAS.

Furthermore, using FC solutions can get better performance if there is more budget.

2.4. Case3 – Extreme High Throughput

Sometimes, for environments such as video editing / streaming media applications, latency is not a problem at all. The only requirement is the high continuous throughput that storage can

provide. High-density I/O ports of QSAN host cards meet this demand. The onboard 10GbE ports plus additional 32Gb / 16Gb FC ports can provide excellent performance. Now it can provide more than 12,000MB/s performance in a box.

3. CONCLUSION

QSAN storage is an ideal solution for mission-critical data centers, high-performance computing, virtualization integration, or media and entertainment environment applications. Using QSAN storage to deploy SMB virtualization infrastructure does not require complicated settings. Enjoy stability and high performance and access QSAN's solutions to optimize existing IT resources within the organization. In addition, installation and management can be completed in just a few minutes, saving IT staff a lot of time and cost.

4. APPENDIX

4.1. Apply To

- XEVO firmware 2.0.3 and later
- SANOS firmware 2.0.1a and later
- QSM firmware 3.3.0 and later

4.2. Reference

Software Manuals

- [XEVO Software Manual](#)
- [SANOS Software Manual](#)
- [QSM Software Manual](#)

White Papers

- [How to Enable MPIO in Windows 2008](#)
- [Implement iSCSI Multipath in RHEL 6.5](#)
- [How to Configure iSCSI Initiator in ESX 6.x](#)

Video Tutorials

- [QSAN FAS Tutorials](#)
- [QSAN SAN Tutorials](#)
- [QSAN NXT Tutorials](#)
- [QSAN NAS Tutorials](#)
- [XCubeFAS Best Practice](#)
- [XCubeSAN Best Practice](#)

- [XCubeNAS Best Practice](#)