

Performance Tuning

XCubeNXT Series White Paper

July 2020

ANNOUNCEMENT

Copyright

© 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.

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, XCubeFAS, XCubeNAS, XCubeNXT, and QSAN.com 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 name or their products are the property of their respective owners.

TABLE OF CONTENTS

Announcement.....	2
Notices.....	4
Preface.....	5
Executive Summary.....	5
Audience	5
Technical Support	5
Information, Tip, and Caution.....	5
1. Hybrid Unified Storage Overview	1
2. Storage Concepts and Configurations	2
2.1. Memory Cache Protection Function	2
2.2. Guidelines for Configuration Planning	3
2.3. Configure Storage	6
3. Configure iSCSI.....	8
3.1. Network Configuration	8
3.2. Configure iSCSI in Windows	9
3.3. Configure iSCSI in Linux.....	11
3.4. Configure iSCSI in VMware ESXi.....	13
4. Configure CIFS.....	18
4.1. Network Configuration	18
4.2. Configure CIFS in Windows	19
4.3. Configure CIFS in Linux.....	19
5. Configure NFS	20
5.1. Network Configuration	20
5.2. Configure NFS in Linux	21
5.3. Configure NFS in VMware ESXi	21
6. Conclusion	22
7. References.....	23

NOTICES

This document is applicable to the following XCubeNXT models:

XCubeNXT Storage System 4U 19" Rack Mount Models

MODEL NAME	CONTROLLER TYPE	FORM FACTOR, BAY COUNT, RACK UNIT
XN8024D	Dual Controller	LFF 24-disk 4U Chassis

XCubeNXT Storage System 3U 19" Rack Mount Models

MODEL NAME	CONTROLLER TYPE	FORM FACTOR, BAY COUNT, RACK UNIT
XN8016D	Dual Controller	LFF 16-disk 3U Chassis

Information contained in 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 establishing iSCSI, CIFS, and NFS connections with QSAN XCubeNXT series products in Windows, Linux, and VMware environments. Regardless of the hardware or software parameters in the environment, it is recommended to use the correct configuration to meet the corresponding requirements.

Audience

This document is applicable for QSAN customers and partners who are interested in learning iSCSI, CIFS, or NFS 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. HYBRID UNIFIED STORAGE OVERVIEW

QSAN XCubeNXT series is QSAN next generation enterprise-grade midrange unified storage providing various high-end enterprise functions, dual active controller design and up to 99.9999% availability.

XCubeNXT series is built-in 10GbE network port and Hybrid SSD cache. With fully modular design, the main components can be easily replaced and upgraded. Unified storage supports advanced enterprise technology and comprehensive data security, which is the best large-scale storage and multi-function data center for small and medium-sized enterprises.



Figure 1-1 XCubeNXT Product Series

For the target market of the XCubeNXT series, they can be on the following applications.

1. Virtualization: Certified by the latest VMware, Hyper-V, and the latest Citrix XenServer. Also ideal for virtualization desktop infrastructure (VDI). Allowing hypervisors to provision data intelligently and run even more VMs with a lower total cost of ownership.
2. Enterprise private cloud: With seamless file access across different operating system, Windows AD/LDAP domain integration, and permission control. XCubeNXT series is a reliable solution to build private cloud and hybrid cloud for business.
3. Backup: XCubeNXT series bring built-in features that suit various recovery plans. These services provide backup and disaster recovery plans within a reasonable budget to meet business needs.

This document will introduce storage design concepts and provide direct guidance for customers using QSAN XCubeNXT storage systems in enterprise environments. The correct settings will make the application even more powerful.

2. STORAGE CONCEPTS AND CONFIGURATIONS

This chapter illustrates the **Memory Cache Protection** function, guidelines, and configurations for optimal performance.

2.1. Memory Cache Protection Function

The **Memory Cache Protection** function will affect performance. Users have to understand the differences between these settings.

2.1.1. Enable Memory Cache Protection

Enabling the **Memory cache protection** setting uses the flash module as SLog which can protect dirty data in memory cache. It guarantees low RTO (Recovery Time Objective) and keeps data safe. In default, CTRL1 is master and CTRL2 is slave, and all pools are imported to the master controller.




Figure 2-1 CTRL1 Failed when Enabling Memory Cache Protection

In the initial state, all pools are running on the master controller. If the CTRL1 (master formerly) fails, the pools will fail over to the CTRL2 (master now) and continue to work. It is similar to the process of CTRL2 failure. The conditions of enabling the **Memory cache protection** are on the following.

- The flash module is inserted and operating normally.
- The **Memory cache protection** setting is enabled.



INFORMATION

Enabling the **Memory cache protection** function will ensure data integrity, but sacrifice some performance.

2.1.2. Disable Memory Cache Protection

When the **Memory cache protection** setting is disabled, the system does not use the SLog mechanism. If there are two or more pools, they must be evenly distributed to the two controllers. Therefore, since both controllers work simultaneously, performance can be improved.




Figure 2-2 CTRL1 Failed when Disabling Memory Cache Protection

In the initial state, pools are averagely imported into dual controller. If the CTRL1 fails, the Pool1 will fail over to the CTRL2. The conditions of disabling the **Memory cache protection** are on the following.

- The **Memory cache protection** setting is disabled.
- Regardless of whether the flash memory module is active.



CAUTION

Disabling the **Memory cache protection** function will improve performance, but risk the short tolerance of RPO (Recovery Point Objective).

2.2. Guidelines for Configuration Planning

Before configuring the storage, follow some guidelines to get the best performance.

1. **Disable Memory cache protection:** According to the Section 2.1 – [Memory Cache Protection Function](#), if you can tolerate a short RPO (Recovery Point Objective), disabling **Memory cache protection** setting will increase IOPS by 2 times and write throughput by up to 5 times. Please go Control Panel -> General Settings -> System, uncheck **Memory cache protection** checkbox, and then click the **Apply** button.

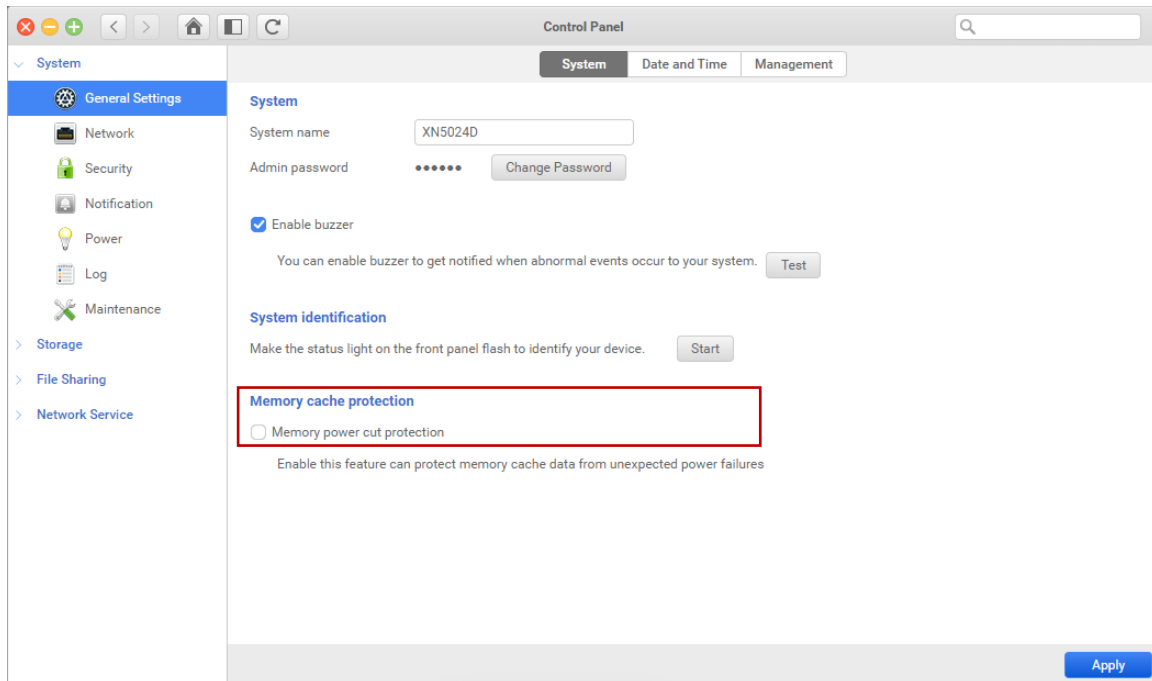


Figure 2-3 Disable Memory Cache Protection



CAUTION

Disabling the **Memory cache protection** function will improve performance, but risk the short tolerance of RPO (Recovery Point Objective). Default is enabled.



TIP

It is recommended to connect a UPS (uninterruptible power supply) or generator to prevent power outages.

2. **Create two pools:** The XCubeNXT series features Dual-Active controller architecture. Both controllers concurrently provide storage services in real time. Active-Active architecture doubles the available host bandwidth and cache-hit ratio which ensures there is no wasted resource in the system. Therefore, the best practice is to create at least two pools assigned to each controller, and let both controllers work together.
3. **Create a RAID 50 pool if disks > 8:** If the number of member disks is larger than 8, it is recommended to create a RAID 50 pool for best performance. Otherwise, create a RAID 5 for optimal.

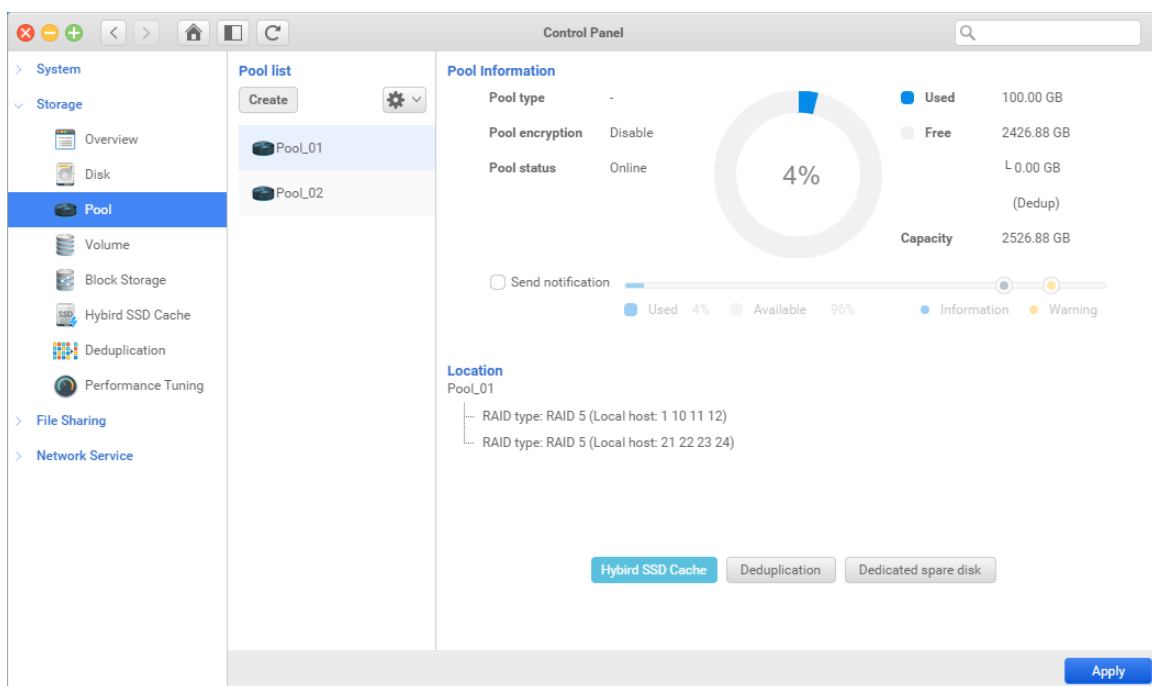


Figure 2-4 Create Two RAID 50 Pools

4. **Enable hybrid SSD cache:** Except uses all flash (all SSDs), putting some SSDs in a traditional HDD pool and enabling hybrid SSD cache can improve performance.

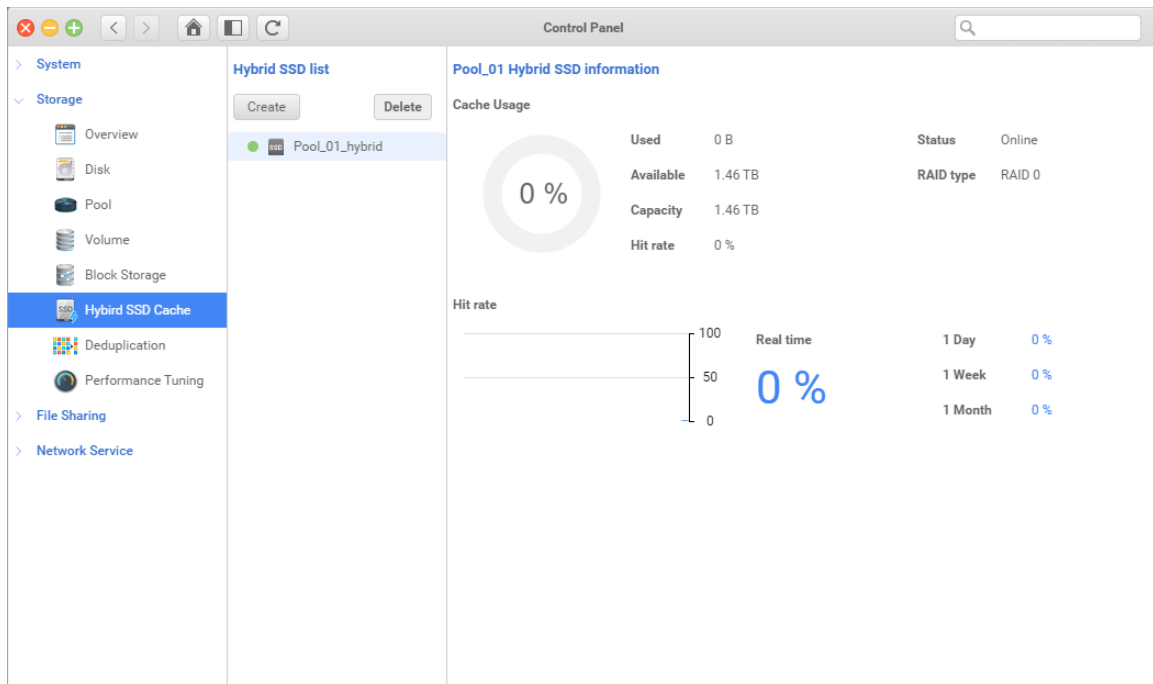


Figure 2-5 Enable hybrid SSD cache

2.3. Configure Storage

Now you have the entire plan to configure storage. The following are the steps to configure storage settings.

1. Follow this video clip - [QSAN NAS Tutorial: NAS Storage Overview and connecting CIFS from Windows](#) to create a pool and a volume.

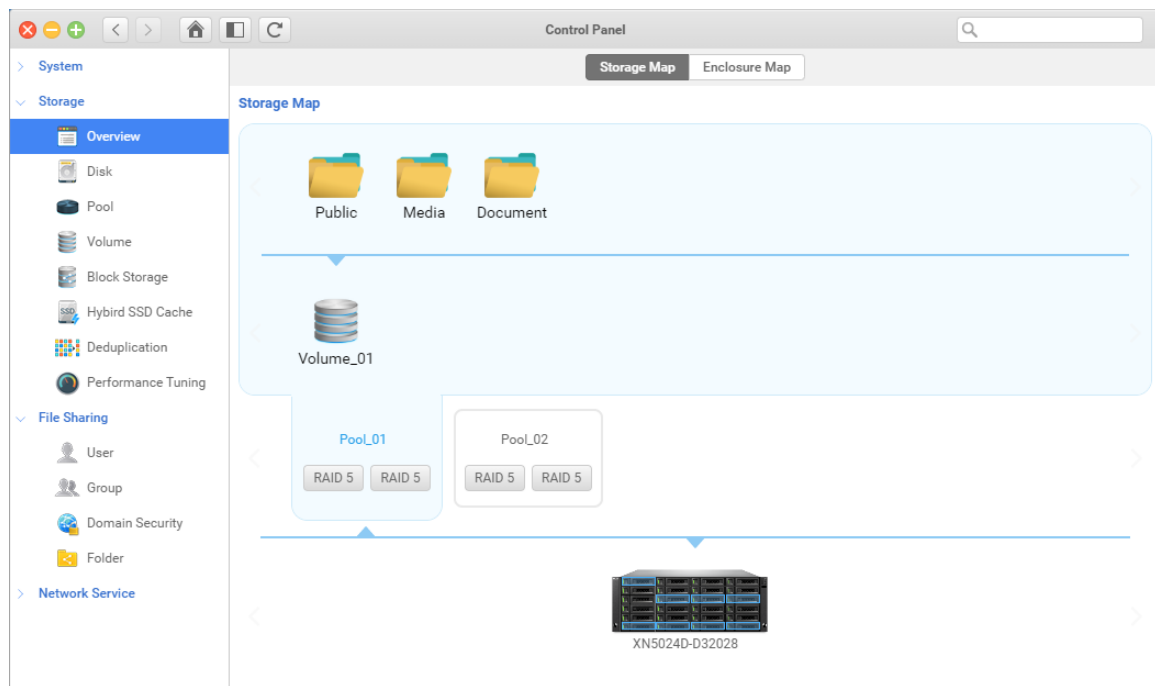


Figure 2-6 Configure Storage

3. CONFIGURE iSCSI

This chapter illustrates the configuration for iSCSI in Microsoft Windows Server, Linux, and VMware ESXi server.

3.1. Network Configuration

Make sure that all the IP addresses of the NIC ports to be used for iSCSI connection can ping to the LAN ports on the XCubeNXT system, please consider configuring different network segments for each NIC port and LAN port, this will make it easier to distinguish and troubleshoot afterward. If there are 4 NIC ports on the server side, the example of the configurations are on the following.

- Server1 - NIC1 (192.168.1.1/24) -->> XCubeNXT - CTRL1-LAN1 (192.168.2.1/24)
- Server1 - NIC2 (192.168.2.1/24) -->> XCubeNXT - CTRL2-LAN2 (192.168.2.2/24)
- Server2 - NIC1 (192.168.3.1/24) -->> XCubeNXT - CTRL1-LAN1 (192.168.3.2/24)
- Server2 - NIC2 (192.168.4.1/24) -->> XCubeNXT - CTRL2-LAN2 (192.168.4.2/24)




Figure 3-1 Topology for iSCSI

For the network switch device, the settings of each switch vendor may be different although the concepts are the same. We usually check the following options.

1. **Jumbo Frame:** Depending on the configuration in the environment, then set the same number on the NIC port.
2. **Flow Control:** Sometimes it must be set to ON. But in some cases we encounter, it is indeed necessary to set it to OFF. Please check the environment and verify performance after changing this setting.
3. **Trunking / LACP:** The existing network configuration in the environment must be implemented first. If LACP was previously enabled, consider using the same configuration. As described earlier in this section, unless you connect the storage to a multi-client topology (ex.: up to 10 or more hosts), it is not recommended to setup Trunking or LACP, although doing so can simplify the connection. Enabling MPIO on the connected host side will be suitable and sufficient.



TIP

After completing all the above suggestions, if you still experience performance degradation, please enable **port mirroring** on the switch for debugging. Connect the network cable from the Windows or MacOS client to the mirror port on the switch, and then install [Wireshark](#) software on the client. Capturing network packets during performance testing sharing the results with the [QSAN technical support team](#), you will get help to find the cause of performance degradation.

3.2. Configure iSCSI in Windows

The following are the steps to configure iSCSI in Microsoft Windows server.

1. Follow this video clip - [Quick Installation setup with XCubeNXT\(Cluster CIFS/NFS, MPIO iSCSI LUN\)](#) to log in to the iSCSI target and configure MPIO from the server side. Remember to reboot the server after the MPIO configuration is complete.
2. If you experience performance degradation when testing performance, try adjusting the following parameters on the HBA driver of the server.
 - Access the **Device Manager** in the **Network Adapter**, right-click the NIC port used for performance testing, select the **Properties**, go to the **Advanced**, find the RSS queue (terms of different brands may vary) , and then adjust the **RSS queue** from 8 to 2.

- At the same location as above, find the **Receive Buffer** and adjust the value to the maximum.
- In the same position as above, find the **Transmit Buffers** and adjust the value to the maximum.
- In the same location as above, find the **Interrupt Moderation Rate** and set its value to **Off**.
- If the above adjustments do not help, please open a command line and go further the following options.

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

or

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



INFORMATION

After completing the above parameters, you must log in to the connected iSCSI session again. If the server does not seem to work properly after logging in, you must restart the server.



TIP

After making one of the above adjustments, please verify the performance immediately to find out which adjustment items are suitable for you. There is not necessary to adjust each item.

3.3. Configure iSCSI in Linux

The following are the steps to configure iSCSI in Linux.

1. For the network configuration in Linux, it is recommended to set the IP addresses to the different network segment.



INFORMATION

It is unable to specify the source NIC port to log in to the target iSCSI portal defined on Linux in a standard manner. If all IP addresses are in the same network segment, the operating system will always use the same NIC as the source to login to a different iSCSI portal, the overall performance of the source NIC will be limited.

-
2. Follow this video clip - [Quick Installation setup with XCubeNXT\(Cluster CIFS/NFS, MPIO iSCSI LUN\)](#) and [QSAN XCubeNAS Tutorial - How to connect iSCSI target from Windows and Unix-Like OS?](#) to log in to the iSCSI target and complete the login process. Watch this video clip - [XCubeFAS - How to setup Host Configuration & Setup Linux iSCSI initiator and MPIO](#) to learn and configure MPIO from the server side.



TIP

Recommend to set the `rr_min_io` value to "1" in the `multipath.conf` file.

-
3. If you experience performance degradation when testing performance, try adjusting the following parameters on the server.
 - Please refer this white paper - [Implement iSCSI Multipath in RHEL6.5](#) and complete all configurations mentioned in the document.
 - Adjust the **RA (Read Ahead) Buffer** for the RAID volume through the `blockdev` command. Set the value of each RAID volume to a higher value, such as 4096 or 8192. If the RAID volume is a multipath device, adjust the `dm-x` device as well.

```
root# 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      1071644682  /dev/sda5
root# blockdev --setra 4096 /dev/sda
```

- Adjust the **Receive Buffer** to 524284 or higher (twice the value) for TCP/IP.

```
root# sysctl -w net.core.rmem_max=524284
net.core.rmem_max = 524284
```

- Disable HT (Hyper Threading) in BIOS. Select the **Advanced** -> **CPU Configuration** -> **Intel (R) HT Technology** -> Change the option to **Disabled**.



INFORMATION

After completing the above parameters, you must log in to the connected iSCSI session again. If the server does not seem to work properly after logging in, you must restart the server.



TIP

After making one of the above adjustments, please verify the performance immediately to find out which adjustment items are suitable for you. There is not necessary to adjust each item.

3.4. Configure iSCSI in VMware ESXi

The following are the steps to configure iSCSI in VMware ESXi server.

1. For the network configuration in VMware ESXi server, it is recommended to set the IP addresses to the different network segment.



INFORMATION

It is unable to specify the source NIC port to log in to the target iSCSI portal defined on Linux in a standard manner. If all IP addresses are in the same network segment, the operating system will always use the same NIC as the source to login to a different iSCSI portal, the overall performance of the source NIC will be limited.

-
2. Follow this video clip - [Quick Installation setup with XCubeNXT\(Cluster CIFS/NFS, MPIO iSCSI LUN\)](#) and [Virtualizing Your iSCSI LUN in VMware with the Simplest Way](#) to log in to the iSCSI target and complete the login process. Read this white paper - [How to configure iSCSI initiator in ESXi 6.x](#) page 18 to learn and configure MPIO from the ESXi server side.



INFORMATION

Please DO NOT create a 4K block size volume if you are going to use this volume in VMware ESXi environment, Because VMware does not yet support 4K block size external storage (Feb. 2020).



TIP

It is recommended to set the PSP (path selection plug-in) to Round Robin, and remember to use VMware Knowledge Base - [Adjusting Round Robin IOPS limit from default 1000 to 1 \(2069356\)](#) to adjust the IOPS value from 1000 to 1.

3. If you want to add another HBA (Host Bus Adapter), please make sure to follow the instructions below for proper configuration. Assign the PCIE device in the ESXi server to the VM.

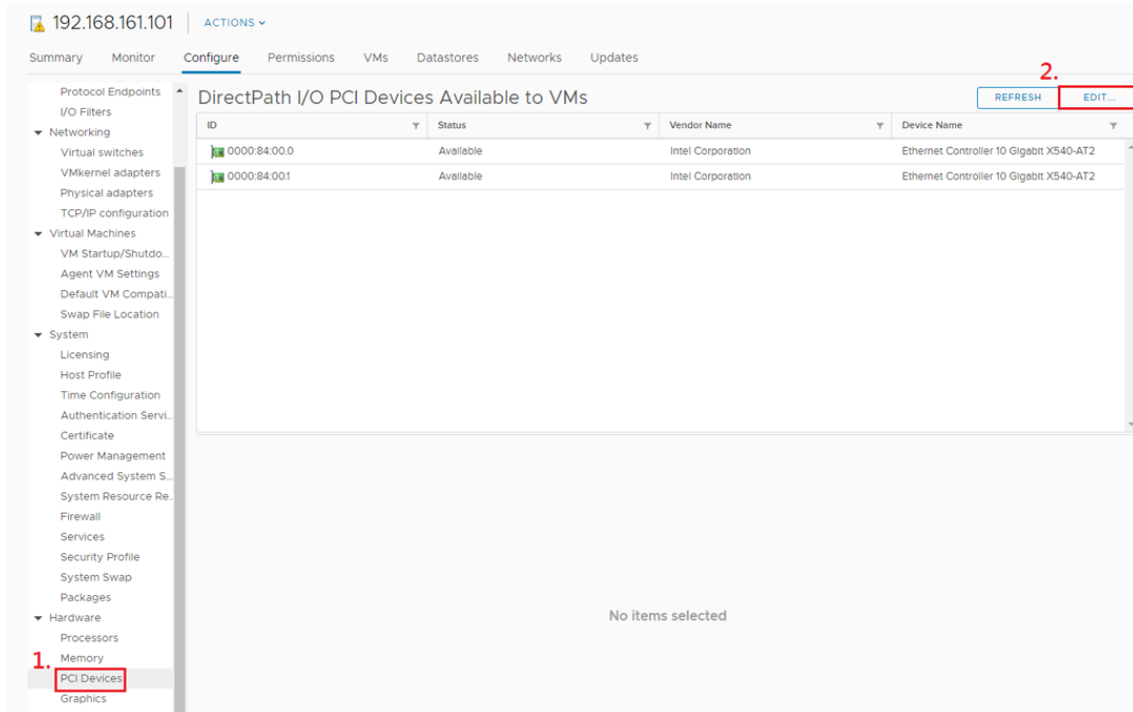


Figure 3-2 Add an HBA in ESXi Server

4. Check the PCI devices. After performing this operation, you have to reboot the ESXi server.

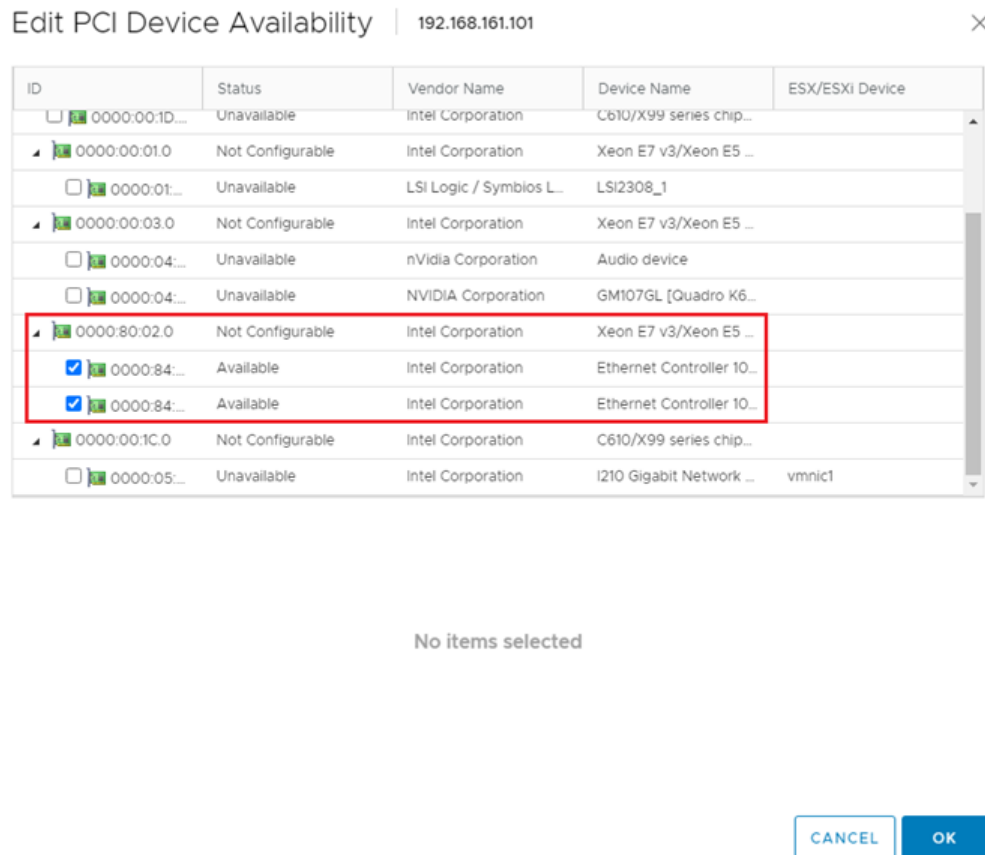


Figure 3-3 Check the PCI Devices

5. Visit the VM and edit the device.

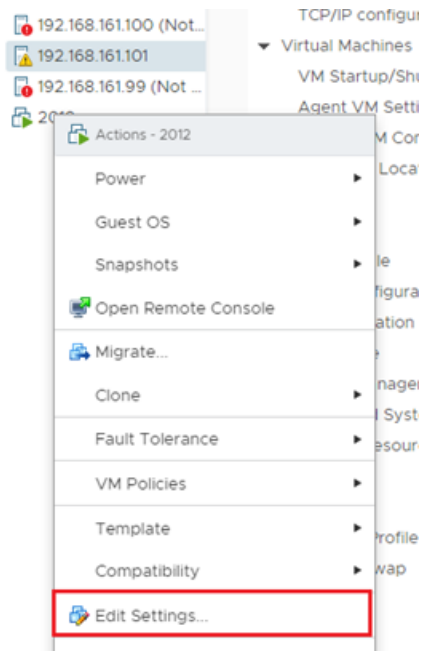


Figure 3-4 Edit the Device

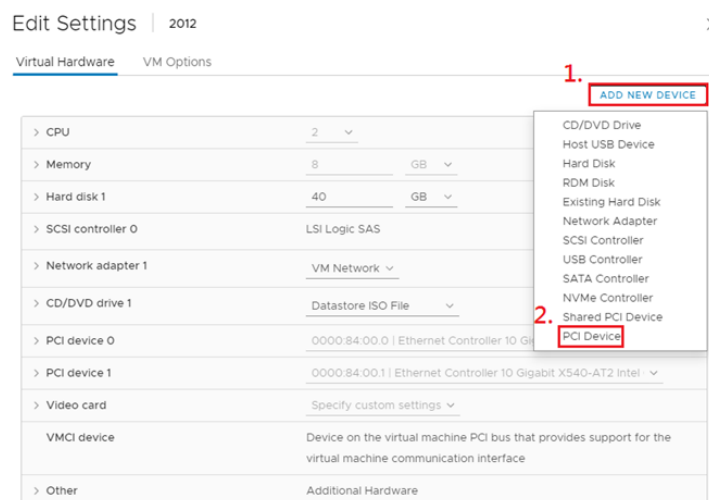


Figure 3-5 Add New PCI Device

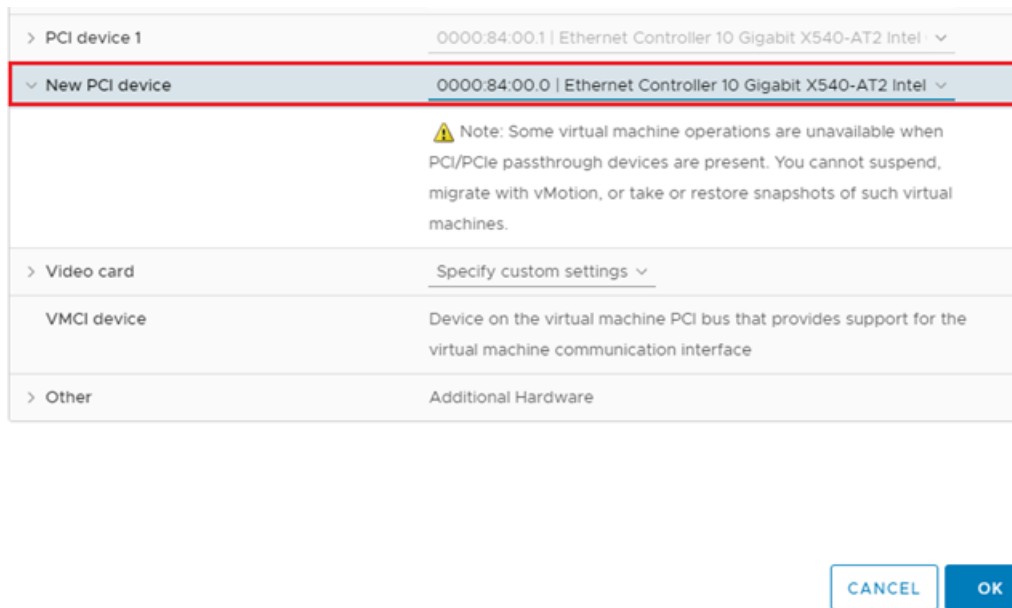


Figure 3-6 New PCI Device is Added

**TIP**

It is recommended add the HBA into VMware. Otherwise, performance may drop a lot if using the default settings.

**INFORMATION**

If you encounter any issues regarding the disconnection between the VMware ESXi server and the storage, VMware Support may ask you the question about the ACK delay function. QSAN products support delayed ACK. The storage side is the receiver of this function. We have our own mechanism to control the queue buffer, and then ACK to the initiator (client) for optimizing performance. In fact, disabling ACK on the ESXi server will not affect the iSCSI service at all. If this can help customers and make them satisfied, it can be adjusted on the ESXi side.

4. CONFIGURE CIFS

This chapter illustrates the configuration for CIFS in Microsoft Windows Server and Linux.

4.1. Network Configuration

For CIFS usage, it has to setup a cluster IP for binding two private IPs. For example:

↓ XCubeNXT - CTRL1-LAN1 (10.10.1.21/24)

- Server - NIC1 (192.168.1.1/24) -->> XCubeNXT - Cluster IP1 (192.168.1.2/24)

↑ XCubeNXT - CTRL2-LAN1 (10.10.1.22/24)

↓ XCubeNXT - CTRL1-LAN2 (10.10.2.21/24)

- Server - NIC2 (192.168.2.1/24) -->> XCubeNXT - Cluster IP2-LAN2 (192.168.2.2/24)

↑ XCubeNXT - CTRL2-LAN2 (10.10.2.22/24)

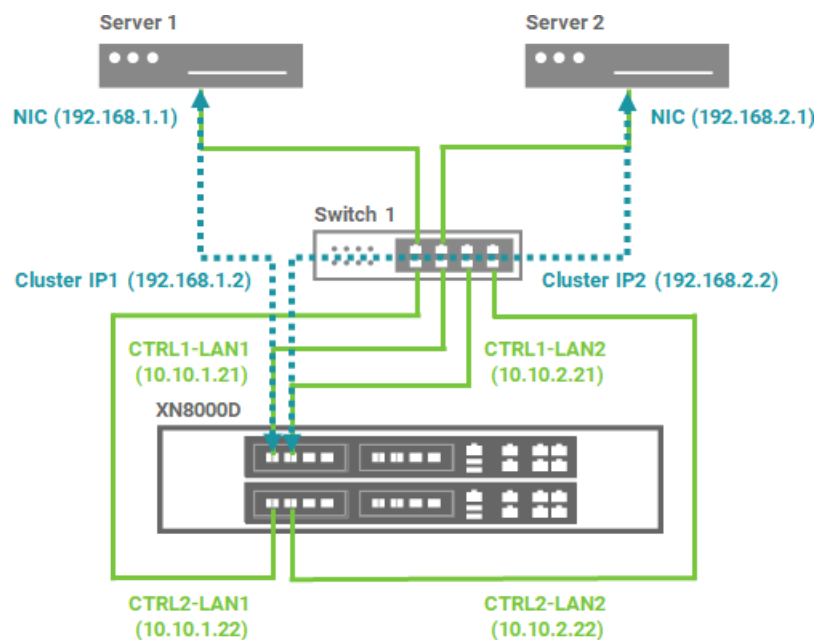


Figure 4-1 Topology for CIFS



INFORMATION

Cluster IP in dual controller will work once at a time.



INFORMATION

IPv6 does not support cluster IP yet, connect to the private IPv6 directly.
(Jun. 2020)

4.2. Configure CIFS in Windows

The following are the steps to configure CIFS in Microsoft Windows server.

1. Follow this video clip - [Quick Installation setup with XCubeNXT\(Cluster CIFS/NFS, MPIO iSCSI LUN\)](#) to configure share folders and complete the login process.
2. If you experience performance degradation when testing performance, please refer to the Section 3.2. - [Configure iSCSI in Windows](#) and try adjusting the parameters on the HBA driver of the server.

4.3. Configure CIFS in Linux

The following are the steps to configure CIFS in Linux.

1. Follow this video clip - [Quick Installation setup with XCubeNXT\(Cluster CIFS/NFS, MPIO iSCSI LUN\)](#) and [QSAN NAS Tutorial: Connect CIFS Shared Folder - Linux](#) to configure share folders and complete the login process.
2. If you experience performance degradation when testing performance, please refer to the Section 3.3. - [Configure iSCSI in Linux](#) and try adjusting the parameters on the server.

5. CONFIGURE NFS

This chapter illustrates the configuration for NFS in VMware ESXi server.

5.1. Network Configuration

Make sure that all the IP addresses of the NIC ports to be used for NFS connection can ping to the LAN ports on the XCubeNXT system, please consider configuring different network segments for each NIC port and LAN port, this will make it easier to distinguish and troubleshoot afterward. If there are 4 NIC ports on the server side, the example of the configurations are on the following.

- Server1 - NIC1 (192.168.1.1/24) --> XCubeNXT - CTRL1-LAN1 (192.168.2.1/24)
- Server1 - NIC2 (192.168.2.1/24) --> XCubeNXT - CTRL2-LAN2 (192.168.2.2/24)
- Server2 - NIC1 (192.168.3.1/24) --> XCubeNXT - CTRL1-LAN1 (192.168.3.2/24)
- Server2 - NIC2 (192.168.4.1/24) --> XCubeNXT - CTRL2-LAN2 (192.168.4.2/24)




Figure 5-1 Topology for NFS

5.2. Configure NFS in Linux

The following are the steps to configure NFS in Linux.

1. Follow this video clip - [Quick Installation setup with XCubeNXT\(Cluster CIFS/NFS, MPIO iSCSI LUN\)](#) and [XCubeNAS - Change Permission from Linux OS after Mounting NFS Shared Folder](#) to mount NFS share folders.
2. If you experience performance degradation when testing performance, please refer to the Section 3.3. - [Configure iSCSI in Linux](#) and try adjusting the parameters on the server.

5.3. Configure NFS in VMware ESXi

The following are the steps to configure NFS in VMware ESXi server.

1. Follow this video clip - [Quick Installation setup with XCubeNXT\(Cluster CIFS/NFS, MPIO iSCSI LUN\)](#) and [QSAN XCubeNAS Tutorial - Mounting NFS Shared Folder from VMware vSphere](#) to mount NFS share folders.
2. If you experience performance degradation when testing performance, please refer to the Section 3.4. - [Configure iSCSI in VMware ESXi](#) and follow the instructions.

6. CONCLUSION

This document provides storage concepts and direct guidance for customers using QSAN XCubeNXT storage systems in enterprise environments. The focus is on system performance and maximizing the ease of use of the dual-controller storage features, giving guidance to achieve optimal performance.

7. REFERENCES

There are some related materials for references.

Video Clips

- [Quick Installation setup with XCubeNXT\(Cluster CIFS/NFS, MPIO iSCSI LUN\)](#)
- [QSAN NAS Tutorial: NAS Storage Overview and connecting CIFS from Windows](#)
- [QSAN XCubeNAS Tutorial - How to connect iSCSI target from Windows and Unix-Like OS?](#)
- [QSAN SAN Tutorial - Setup Windows iSCSI Initiator with MPIO](#)
- [XCubeFAS - How to setup Host Configuration & Setup Linux iSCSI initiator and MPIO](#)
- [Virtualizing Your iSCSI LUN in VMware with the Simplest Way](#)
- [QSAN NAS Tutorial: Connect CIFS Shared Folder - Linux](#)
- [XCubeNAS - Change Permission from Linux OS after Mounting NFS Shared Folder](#)
- [QSAN XCubeNAS Tutorial - Mounting NFS Shared Folder from VMware vSphere](#)

Documents

- Best Practice - [iSCSI Performance Tuning](#)
- Best Practice - [How to Adjust Performance in Windows](#)
- Best Practice - [How to Adjust Performance in Linux](#)
- Best Practice - [How to Adjust Performance in VMware](#)
- White Paper - [Implement iSCSI Multipath in RHEL6.5](#)
- White Paper - [How to configure iSCSI initiator in ESXi 6.x](#)
- VMware Knowledge Base - [Adjusting Round Robin IOPS limit from default 1000 to 1 \(2069356\)](#)