

How to Adjust Performance in Linux

XCubeNXT Series Best Practice

July 2020

OVERVIEW

This document will guide customers how to use the XCubeNXT storage system to adjust performance in Linux. The correct settings will make the application more powerful. The applicable models are XN8024D and XN8016D.

CONFIGURE STORAGE

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

1. **Disable Memory cache protection:** 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.



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.

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.

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.

CONFIGURE ISCSI IN LINUX

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)

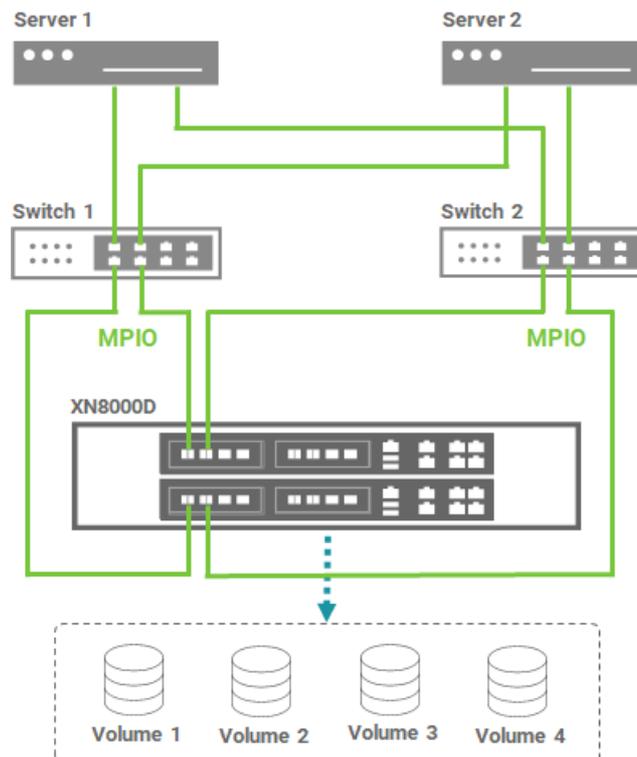


Figure 1 Topology for iSCSI

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 - [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.

CONFIGURE CIFS IN LINUX

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 (192.168.2.2/24)

↑ XCubeNXT - CTRL2-LAN2 (10.10.2.22/24)

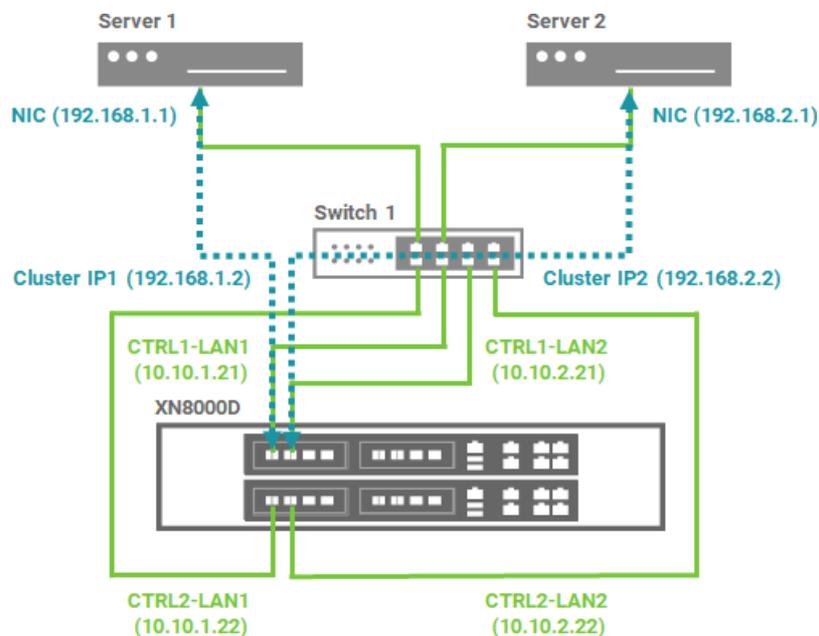


Figure 0-1 Topology for CIFS



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)

The following are the steps to configure CIFS in Linux.

1. Follow this video clip - [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 - [Configure iSCSI in Linux](#) and try adjusting the parameters on the server.

CONFIGURE NFS IN LINUX

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)

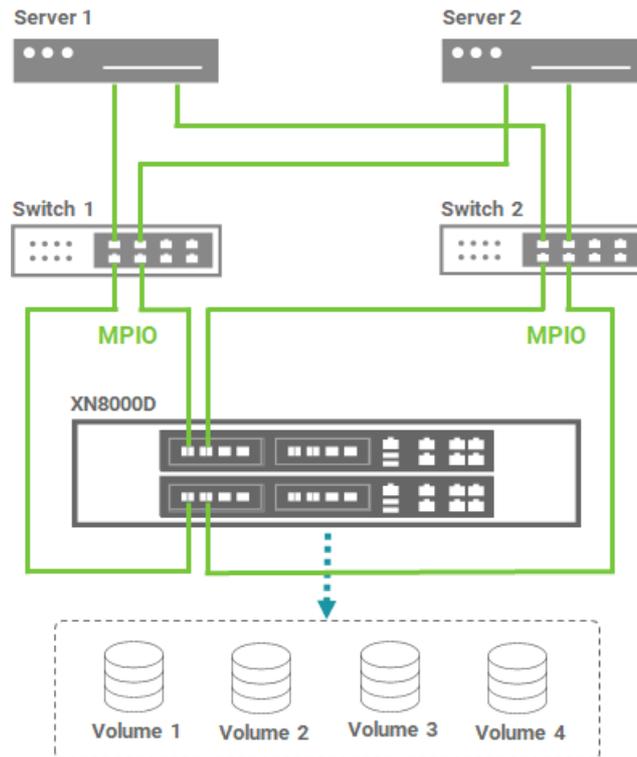


Figure 2 Topology for NFS

The following are the steps to configure NFS in Linux.

1. Follow this video clip - [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 - [Configure iSCSI in Linux](#) and try adjusting the parameters on the server.

REFERENCES

There are some related materials for references.

Video Clips

- [QSAN NAS Tutorial: NAS Storage Overview and connecting CIFS from Windows](#)
- [XCubeFAS - How to setup Host Configuration & Setup Linux iSCSI initiator and MPIO](#)
- [QSAN NAS Tutorial: Connect CIFS Shared Folder - Linux](#)
- [XCubeNAS - Change Permission from Linux OS after Mounting NFS Shared Folder](#)

Documents

- Best Practice - [iSCSI Performance Tuning](#)

- White Paper - [Implement iSCSI Multipath in RHEL6.5](#)
- White Paper - [XCubeNXT Performance Tuning](#)