



# BEAST and E-Series

## Multipathing Best Practices Guide

## Trademarks

Nexsan<sup>®</sup>, BEAST<sup>™</sup>, BEAST P<sup>™</sup>, BEAST Elite<sup>™</sup>, BEAST X<sup>™</sup>, Nexsan E60<sup>™</sup>, Nexsan E60V<sup>™</sup>, Nexsan E60VT<sup>™</sup>, Nexsan E60P<sup>™</sup>, Nexsan E60X<sup>™</sup>, Nexsan E48<sup>™</sup>, Nexsan E48V<sup>™</sup>, Nexsan E48P<sup>™</sup>, Nexsan E48VT<sup>™</sup>, Nexsan E48X<sup>™</sup>, Nexsan E32F<sup>™</sup>, Nexsan E32V<sup>™</sup>, Nexsan E32X<sup>™</sup>, Nexsan E18<sup>™</sup>, Nexsan E18P<sup>™</sup>, Nexsan E18F<sup>™</sup>, Nexsan E18V<sup>™</sup>, Nexsan E18X<sup>™</sup>, and the Nexsan logo are trademarks or registered trademarks of Nexsan.

All other trademarks and registered trademarks are the property of their respective owners.

## Patents

This product is protected by one or more of the following patents, and other pending patent applications worldwide:

United States patents US8,191,841, US8,120,922

United Kingdom patents GB2296798B, GB2297636B, GB2466535B, GB2467622B, GB2467404B

## Regulatory compliance

United States Statement for FCC: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his own expense.

Electromagnetic Emissions: FCC Class A, EN 55022 Class A, EN 61000-3-2/-3-3, CISPR 22 Class A

Electromagnetic Immunity: EN 55024/CISPR 24, (EN 61000-4-2, EN 61000-4-3, EN 61000-4-4, EN 61000-4-5, EN 61000-4-6, EN 61000-4-8, EN 61000-4-11)

Safety: CSA/EN/IEC/UL 60950-1 Compliant, UL or CSA Listed (USA and Canada), CE Marking (Europe)

California Best Management Practices Regulations for Perchlorate Materials: This Perchlorate warning applies only to products containing CR (Manganese Dioxide) Lithium coin cells. Perchlorate Material-special handling may apply. See [www.dtsc.ca.gov/hazardouswaste/perchlorate](http://www.dtsc.ca.gov/hazardouswaste/perchlorate).

## Use and limitations of this document

Unauthorized use, duplication, or modification of this document in whole or in part without the written consent of Nexsan is strictly prohibited.

Nexsan reserves the right to make changes to this manual, as well as the equipment and software described in this manual, at any time without notice. This manual may contain links to Web sites that were current at the time of publication, but have since been moved or become inactive. It may also contain links to sites owned and operated by third parties. Nexsan is not responsible for the content of any such third-party site.

# Contents

---

- About this manual** ..... **v**
  - Conventions ..... v
  - Notes, tips, cautions, and warnings ..... v
  - Contacting Nexsan ..... vi
    - Service and support ..... vi
  
- Chapter 1: Introduction to multipathing** ..... **9**
  - Overview ..... 9
  - How multipathing works ..... 10
  - When multipathing is required ..... 10
  - Why multipathing is useful ..... 11
    - Fault tolerance ..... 11
    - High availability ..... 11
    - Improved performance ..... 11
  - Volume mapping, system modes, and multipathing ..... 11
  - Load balancing policies ..... 12
  - Recommended Nexsan Storage System settings ..... 13
  
- Chapter 2: Multipathing on Windows Server** ..... **15**
  - Prerequisites ..... 15
  - Installing Microsoft Multipathing (Windows Server 2016 and later) ..... 15
  - Installing Nexsan MultiSan IO (Windows Server 2012 R2 and earlier) ..... 16
  - Setting up multipathing ..... 18
    - Verifying that multipathing is enabled ..... 18
    - Enabling the iSCSI Initiator service ..... 21
    - Connecting to a target storage device over iSCSI ..... 22
  - Recommended disk timeout settings ..... 23
  
- Chapter 3: Multipathing on Linux** ..... **25**
  - Prerequisites ..... 25
  - Multipathing components and definitions ..... 25
    - multipath.conf ..... 26
    - Attributes and Values ..... 27
  - Configuring Linux multipathing ..... 28

Multipath.conf examples .....	31
Redhat 5/Centos .....	31
Redhat 6 .....	32
DEBIAN Lenny .....	33
DEBIAN Wheezy .....	34
SUSE 11 (including OpenSUSE) .....	36
<b>Chapter 4: Multipathing on Solaris .....</b>	<b>39</b>
Prerequisites .....	39
Multipath components .....	40
Multipathed device GUID .....	40
Configuring Solaris I/O multipathing .....	41
Rolling restarts and configurable timeouts .....	45
SCSI device timeout and retry .....	45
Transport timeouts .....	45
Fibre Channel transport .....	45
iSCSI transport .....	46
<b>Chapter 5: Multipathing on IBM AIX .....</b>	<b>47</b>
IBM AIX Multiple Path IO .....	47
Nexsan Multi-Path I/O for IBM AIX .....	49
Installing MSIO software for IBM AIX .....	49
Discover MSIO volumes .....	50
Viewing and changing MSIO policy .....	50
Volume groups .....	52
HACMP best practices .....	52
Uninstalling MSIO .....	52
<b>Chapter 6: Multipathing on XenServer .....</b>	<b>55</b>
Configuring multipathing on XenServer .....	55
Multipathing.conf example .....	55
<b>Glossary .....</b>	<b>57</b>

## About this manual

This best practices guide provides information for implementing multipathing on Nexsan E-Series and Nexsan BEAST Storage Systems.

**Note** While Nexsan makes every effort to ensure the accuracy of technical documentation, screen images and procedures may change after publication. In case of discrepancy, please check for the latest updates on the E-Series and BEAST [Documents and Downloads](#) page. Also, refer to the latest Release Notes.

## Conventions

Here is a list of text conventions used in this document:

Convention	Description
<u>underlined blue</u>	Cross-references, hyperlinks, URLs, and email addresses.
<b>boldface</b>	Labels on the physical Nexsan Storage System or interactive items in the graphical user interface (GUI).
<i>italics</i>	System messages and non-interactive items in the GUI. References to software user guides.
monospace	Command-line interface (CLI) text or text that refers to file or directory names.
<b>monospace bold</b>	Text strings that must be entered by the user in the CLI or in text fields in the GUI.

## Notes, tips, cautions, and warnings

**Note** Notes contain important information, present alternative procedures, or call attention to certain items.

**Tip** Tips contain handy information for end-users, such as other ways to perform an action.



**CAUTION:** In hardware manuals, cautions alert the user to items or situations which may cause damage to the Nexsan Storage System or result in mild injury to the user, or both. In software manuals, cautions alerts the user to situations which may cause data corruption or data loss.



**WARNING:** Warnings alert the user to items or situations which may result in severe injury or death to the user.

## Contacting Nexsan

For questions about Nexsan products, please visit the [Nexsan support](#) Web page, and the E-Series and BEAST [Documents and Downloads](#) page. If you are unable to find the answer to your question there, please see our contact information below.

### Service and support

Nexsan's Technical Services Group provides worldwide assistance with installation, configuration, software support, warranty, and repair for all Nexsan products. A variety of service and support programs are available to provide you with the level of coverage and availability your operation requires.

#### **Nexsan Headquarters**

1289 Anvilwood Avenue  
Sunnyvale, CA 94089  
United States of America

#### **Worldwide Web site**

[www.nexsan.com](http://www.nexsan.com)

#### **E-Series and BEAST support:**

[https://helper.nexsansupport.com/esr\\_support](https://helper.nexsansupport.com/esr_support)

#### **European Head Office, UK**

Units 33–35 Parker Centre  
Mansfield Road  
Derby, DE21 4SZ  
United Kingdom

#### **Contact:**

<https://helper.nexsansupport.com/contact>

## Related documents

The following Nexsan product manuals contain related information:

- *Nexsan High-Density Storage User Guide*
- *Nexsan Snapshots and Replication User Guide*
- *Nexsan Storage Tools online help and user guide (bundled with the tools)*

## Revision history

This section lists updates and new material added to the *Nexsan Multipathing Best Practices Guide*.

### P0450148, Rev. C, March 2022

Updated for technical accuracy, applied new Nexsan template and branding.

### P0450148, Rev. B, August-September 2019

Updated the Introduction and Windows Server chapters, specifying the software and procedures required to set up multipathing on supported Windows Server systems.

Updated procedures in the Solaris and Linux chapters.

Updated the VMware chapter and moved the content to the *VMware Best Practices Guide*.

### P0450148, Rev. A, July 2016

Updates for E-Series 18 and E-Series 32.

### NXS-BEMP-BP Rev. 03, November 2015

- Changed name of document to *Nexsan High-Density Storage Multipathing Best Practices Guide*.
- Made changes throughout to reflect the new Nexsan BEAST line of RAID storage systems.

### NXS-LSMP-BP Rev. 02, October 2014

Minor wording changes and corrections throughout; added note to [Installing Nexsan MultiSan IO \(Windows Server 2012 R2 and earlier\)](#) on page 16 about installing the Windows Server multipathing role; changed step 1 in [Installing Nexsan MultiSan IO \(Windows Server 2012 R2 and earlier\)](#) on page 16 to state that Nexsan Storage System and host machine should be on the same subnet; added troubleshooting suggestions to step 1 in [Configuring Linux multipathing](#) on page 28; added Chapter 6, [Multipathing on XenServer](#) on page 55.

### NXS-LSMP-BP Rev. 01, May 2014

Changed formatting throughout to reflect Nexsan as an Nexsan brand.





## Introduction to multipathing

---

This chapter contains the following sections:

Overview .....	9
How multipathing works .....	10
When multipathing is required .....	10
Why multipathing is useful .....	11
Volume mapping, system modes, and multipathing .....	11
Load balancing policies .....	12
Recommended Nexsan Storage System settings .....	13

### Overview

Multipathing enables I/O commands to route over multiple connections between hosts and Nexsan Storage Systems. When multiple connections are established, some form of multipathing policy is required to manage multiple instances of volumes, which will in turn ensure high availability, fault tolerance, and improved performance.



**CAUTION:** If multipathing is not properly installed and configured, the host may view each instance of a volume as a separate and independent disk device, which could result in accidental disk erasure and data corruption.

This document provides multipathing configuration details for the following operating systems:

- Windows Server: see "Multipathing on Windows Server" ([page 15](#))
- Linux: see "Multipathing on Linux" ([page 25](#))
- Solaris: see "Multipathing on Solaris" ([page 39](#))
- IBM AIX: see "Multipathing on IBM AIX" ([page 47](#))
- XenServer: see "Multipathing on XenServer" ([page 55](#))

To configure multipathing for VMware, please see the *Nexsan VMware Best Practices Guide* on the Nexsan E-Series/BEAST Support [Documents & Downloads](#) page.


## How multipathing works

Normally, a disk or volume has only one path between it and the host accessing it. If this path should fail, then it becomes impossible to access the disk or volume. Multipathing enables a single volume to be connected to a host via two or more paths so that the failure of one path does not make the volume inaccessible.

In this sample configuration (called “no single point of failure”), each volume is available to the host through as many as eight different paths due to the interconnection of the volume, the RAID controllers, and the switches.

**Note** This configuration assumes that the Nexsan Storage System is in All Ports All LUNs (APAL) mode.

Figure 1-1: "No single point of failure" multipathing configuration (example)



Each volume is mapped to both ports on each controller, and each controller has a path to each switch. Without multipathing, the server would see each volume as eight different devices. With multipathing, however, each volume is only presented to the server once, but each path remains active. This configuration can survive the failure of any of the following without losing connectivity to the volumes:

- Either path between the server and the switches
- Either switch, either RAID Controller, or even one of each
- Up to three of the four paths between the switches and the RAID Controllers
- Up to three of the four ports on the RAID Controllers

## When multipathing is required

Multipathing is required wherever a given volume can present more than one path to a host server. Depending on your configuration, any of the following may cause a volume to present more than one path to a host server:

- A storage volume being mapped to more than one host port on the RAID Controllers.
- Multiple paths between RAID Controllers and switches.
- More than one switch.
- Multiple paths between switches and the host server.

## Why multipathing is useful

Although it is possible to set up a storage installation or a storage area network (SAN) so that volumes only present one path to each host server, there are three main reasons to use multipathing: fault tolerance, high availability, and improved performance.

### Fault tolerance

In an installation where volumes present multiple paths to host servers, the failure of any one path doesn't make the volume inaccessible. Path failure can occur when a RAID Controller, switch, interface card, or port fails. With multipathing in place, such faults are transparent to the end user and can be attended to without any volume down-time.

### High availability

In single-path installations, any required maintenance of the system (repair, replacement, or reconfiguration) requires RAID storage to be taken offline while the maintenance is being performed. With multipathing in place, maintenance can be performed without having to take down any volumes, because at least one additional path maintains the connection between the volume and the host server.

### Improved performance

In some cases, multipathing can provide faster throughput between host servers and volumes because of the availability of additional data paths.

## Volume mapping, system modes, and multipathing

Any time a volume in a Nexsan High-Density Storage unit is mapped to more than one active port, multipathing is required. Volumes are often mapped to multiple active ports when the Nexsan system is in 4-Port Active-Active (4PAA) or All Ports All LUNs (APAL) mode. However, it is also possible to map a volume to both ports on a RAID Controller when the unit is in Single Controller or Dual Controller Non-Redundant (DCNR) mode. In all of these cases, multipathing software is required to manage the multiple paths that the volume presents to the host server.

For more information on volume mapping, system modes, and how they affect multipathing, see the *Nexsan Nexsan High-Density Storage User Guide*.

## Load balancing policies

Load balancing policies are sets of instructions for multipathing software to follow to ensure that I/O transfers through host paths are optimally routed and that no one path becomes overloaded with data. Choose the most suitable MPIO load balancing policy based on your network and storage requirements.

For detailed load balancing examples and configuration steps using the Nexsan MultiSan (MSIO), refer to "Configuration" in the Nexsan Storage Tools online help and user guide, which are included with the tools.

Here is a brief description of each policy type:

Policy	Description
Failover only	<p>Allocates a single active path for each target disk, and any additional paths will be passive. When the active path fails, one of the passive paths will become active, and host I/O will continue down this new path. This is the most basic load balancing mode.</p> <p>When failover only load balancing is active for a particular target disk, it is possible to configure the preferred path; this path will always then be chosen as the active path, providing the path is in a working state. This can ensure that the same path for each target disk is not used for all host I/O, or to route I/O through a better-performing path.</p> <p>It is advised for performance reasons to select a path that connects to the controller that owns the array that the target disk or LUN resides on to avoid I/O having to be routed over the inter-controller link.</p>
Round-robin load balancing	<p>Switches all the paths to become active and sends host I/O down all of the available paths, alternating between all the paths in a round robin or circular fashion. For example, if a target disk has 3 paths, the first host I/O will be sent to path 1, the next host I/O to path 2, then path 3, then back to path 1 and so on.</p> <p>With round robin load balancing, depending on the configuration, improved performance may be achieved due to host I/O being spread across all the available paths.</p>
Round-robin load balancing with subset	<p>Operates in the same way as round-robin load balancing, except it is possible to manually configure individual paths as active or passive using the MSIO or MPIO management utility.</p> <p>This option may be useful if, for example, you know that some of the paths are routed through a lower-performance switch, and you therefore wish to make these paths passive.</p>
Least queue depth	<p>Uses all available paths to your target disks, with host I/O sent to the least busy path based on the number of queued host I/Os.</p> <p>The advantage of this load balancing policy is that host I/O is evenly spread across the available paths, which can result in better performance (depending on configuration).</p>
Weighted paths	<p>Enables assigning a numeric value, known as "weight", to each path. Host I/O is sent along the path with the lowest assigned weight; therefore, I/O will only be active on a single path, regardless of the number of active paths.</p> <p>Use weighted paths to set a priority for each path, so that when path failures occur, I/O will be resumed on the next highest weighted path.</p>

## Recommended Nexsan Storage System settings

Use this procedure for guidance about the recommended software configuration for effective multipathing.

► **To configure multipathing software settings:**

1. For systems with Fibre Channel connections, configure the RAID system mode to **APAL – all ports all LUNs** on your Nexsan Storage System. This mode enables four paths to be active to each configured volume.

The APAL RAID controller mode is recommended for use with multipathing. APAL provides the maximum amount of path redundancy: a single volume can be mapped to both fibre ports on controller 0 and both fibre ports on controller 1, for a total of four paths to any given volume. This also protects against controller failure: when one controller fails, only 2 of the 4 paths fail.

2. For systems with iSCSI connections, refer to "Configure host access", in the *Nexsan High Density Storage User Guide*.
3. Restart the Nexsan Storage System.
4. Create the desired number of arrays and volumes on your Nexsan Storage System. Try to divide the arrays between the controllers (where possible), so that controller 0 "owns" the same number of arrays as controller 1, this will give the best performance and ensures both controllers are active. You may want to spend more time considering which arrays to put on which controller so that both controllers share 50 percent of the I/O load.
5. Map all volumes to all 4 host ports, and configure the access rights for all discovered hosts.
6. Use the **Disk Management** tool for your operating system to ensure that all the newly created volumes on your Nexsan Storage System are present. These new disks can now be formatted and used as available storage.
7. For increased performance, you may want to select a different load balancing mode using the Nexsan MultiSan IO Windows or Multipathing (MPIO) applications. See "Configuring the Default Load Balancing Policy" of the Nexsan Storage Tools user guide.

Optionally, once everything is configured and running to your satisfaction, you may want to create deliberate failure scenarios to test that host I/O continues after the failure and that paths return after the failure is reversed. For example, in a fibre channel-based system, you could remove power to one of the fibre switches, remove a fibre cable, and fail a RAID controller, and so on. In each failure case, carefully work out that at least one path to your configured targets remains.



## Multipathing on Windows Server

---

This chapter contains the following sections:

Prerequisites .....	15
Installing Microsoft Multipathing (Windows Server 2016 and later) .....	15
Installing Nexsan MultiSan IO (Windows Server 2012 R2 and earlier) .....	16
Setting up multipathing .....	18
Recommended disk timeout settings .....	23

### Prerequisites

To enable multipathing on Windows Server hosts and Nexsan Storage, you'll need the following hardware and software packages installed on your Windows Server PCs:

- One or more Nexsan Storage Systems, with dual controllers.
- At least one fibre channel or iSCSI switch and all required cables.
- For redundancy, use at least a dual-port HBA.
- The Windows Server software of your choice, including all Microsoft updates and the latest service pack, where applicable.

### Installing Microsoft Multipathing (Windows Server 2016 and later)

Use this procedure for details about installing the Microsoft Multipathing (MPIO) tools.



**CAUTION:** Installing Multipathing (MPIO) may cause storage volumes to be remounted. You must ONLY install MPIO during scheduled system downtime, during which all host I/O to and from the Nexsan Storage System is stopped.


#### **For more information:**

- In the Nexsan Storage Tools, refer to "Multipathing (MPIO)" in the bundled online help or user guide.
- For Microsoft PowerShell cmdlets, refer to the Microsoft documentation at <https://docs.microsoft.com/en-us/powershell/module/mpio/?view=win10-ps>

► **To install Multipathing (MPIO):**

1. On the Windows Server PC that is to be connected to the Nexsan Storage System, navigate to **NexsanTools\_x.x.x.exe** (where x.x.x is the version number) in the **Tools** folder.
2. Double-click the **NexsanTools\_x.x.x.exe** icon to launch the installer.  
The *Welcome* dialog displays.
3. Click **Next** to display the *End-User License Agreement* and review the terms.
4. Click the check box next to **I accept the terms in the License Agreement**, then click **Next**.  
The *Custom Setup* dialog displays.

Figure 2-1: *Custom Setup* dialog



5. Make sure that **Multipathing (MPIO)** is selected.  
**Note** Installation of the other Nexsan Storage Tools is optional.
6. Click **Next** to go to the **Multipathing (MPIO)** dialog.  
This dialog contains the caution statement which informs you that installing Nexsan MultiSan IO may cause volumes to be remounted.
7. Read the caution text, then check the **Acknowledge this warning and continue** check box.
8. Click **Next** to go to the *Ready to Install Multipathing (MPIO)* dialog.
9. Click **Install** to begin the installation process. This may take a few minutes.  
When the installation is complete, the *Completed the Multipathing Setup Wizard* dialog displays.
10. Click **Finish** to close the dialog.
11. If the *Nexsan Storage Tools Setup* dialog appears stating that you must restart your computer, click **Yes**.

## Installing Nexsan MultiSan IO (Windows Server 2012 R2 and earlier)

Use this procedure for details about installing the Nexsan MultiSan IO (MSIO) multipathing tools. For help about the Nexsan Storage Tools, please refer to the Nexsan Storage Tools user guide, included with the tools.

### *MultiSan IO features*

Here is a brief outline of the features of the Nexsan Storage Tools and their applications:



Feature	Application
Monitoring	Provides detailed information about discovered paths, target disks, and connections.
Configuration	Configures load balancing policy and failover strategy, email alerts, and other options.
Command Line Management	Used to manage MultiSan IO from the command line.
Diagnostics and Troubleshooting	Provides tools and information for troubleshooting MultiSan IO.

### Prerequisites

- The Nexsan Storage Tools and bundled documentation, available for download from the Nexsan Support [Documents & Downloads](#) page.

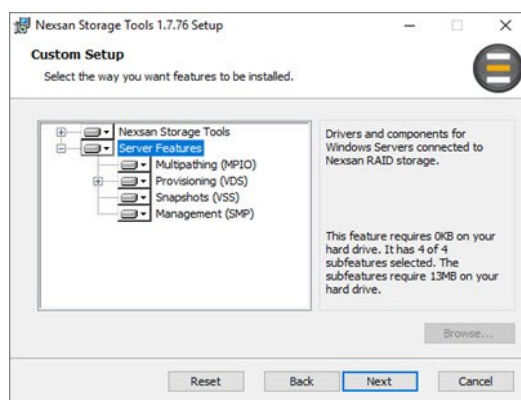


**CAUTION:** Installing Nexsan MultiSan IO may cause storage volumes to be remounted. You must **ONLY** install these tools during scheduled system downtime, during which all host I/O to and from the Nexsan Storage System is stopped.

#### ► To install Nexsan MSIO:

1. On the server that is to be connected to the Nexsan Storage System, navigate to **NexsanTools\_x.x.x.exe** (where x.x.x is the version number) in the **Tools** folder.
2. Double-click the **NexsanTools\_x.x.x.exe** icon to launch the installer.  
The *Welcome* dialog displays.
3. Click **Next** to display the *End-User License Agreement* and review the terms.
4. Click the check box next to **I accept the terms in the License Agreement**, then click **Next**.  
The *Custom Setup* dialog displays.

Figure 2-2: Custom Setup dialog



5. Make sure that **Multipathing** is selected.  
**Note** Installation of the other Nexsan Storage Tools is optional.
6. Click **Next** to open the *Nexsan MultiSan IO Selected* dialog.

This dialog contains the caution statement which informs you that installing Nexsan MultiSan IO may cause volumes to be remounted.

7. Read the caution text, then check the **Acknowledge this warning and continue** check box.
8. Click **Next** to go to the *Ready to Install Nexsan Multipathing* dialog.
9. Click **Install** to begin the installation process. This may take a few minutes.  
When the installation is complete, the *Completed the Nexsan Multipathing Setup Wizard* dialog displays.
10. Click **Finish** to close the dialog.
11. If the *Nexsan Storage Tools Setup* dialog appears stating that you must restart your computer, click **Yes**.

## Setting up multipathing

Use the procedures in this section for setting up multipathing to Nexsan storage devices.

Verifying that multipathing is enabled .....	18
Enabling the iSCSI Initiator service .....	21
Connecting to a target storage device over iSCSI .....	22

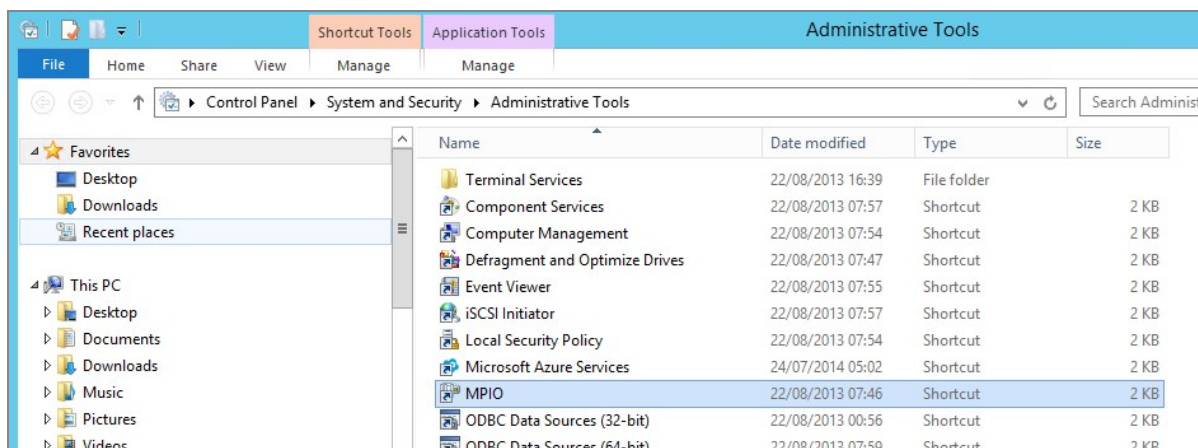
### Verifying that multipathing is enabled

Use the two procedures in this topic to verify that multipathing is enabled for Nexsan Storage Systems.

► **To verify that multipathing is enabled:**

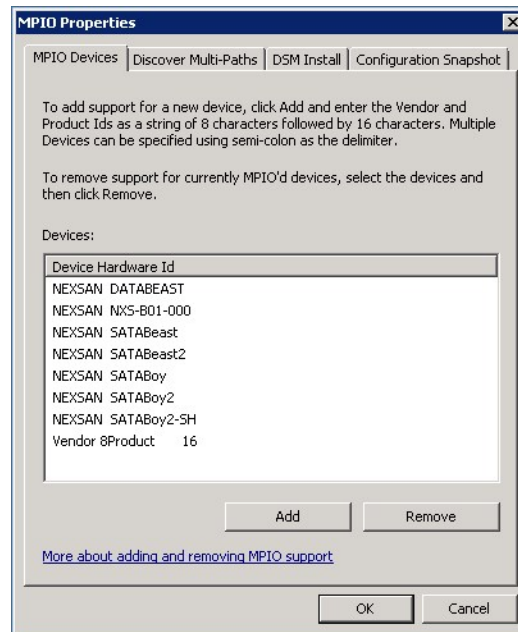
1. Open the *MPIO Properties* dialog:
  - a. In Windows Server 2016 or later, go to **Start > Windows Administrative Tools > MPIO**, or
  - b. In Windows Server 2012 R2 or earlier, the same Microsoft tool is provided under **Start > Administrative Tools > MPIO**.

Figure 2-3: *MPIO Properties* in Administrative Tools



- On the **MPIO Devices** tab, ensure that the Nexsan configurations have been added to the *Devices* list.

Figure 2-4: Devices list in the MPIO Properties dialog



- For Windows 2012 R2 or earlier, do one of the following:
  - For Fibre Channel and SAS connections, connect to the target storage systems using the physical cabling or switch zoning.

Figure 2-5: Nexsan Multipath IO *Target disks* list

Target	Status	LUN	LB Policy	Total Paths	No. Active P
NEXSAN SATABoy2 Volume 1	OK	1	Least queue depth TPGS (D...	1	1
NEXSAN SATABoy2 Volume 2	OK	3	Least queue depth TPGS (D...	1	1

- For iSCSI connections, start the Microsoft iSCSI Initiator Service and make sure that the system is connected to the target storage systems (see "Connecting to a target storage device over iSCSI" (page 22)).

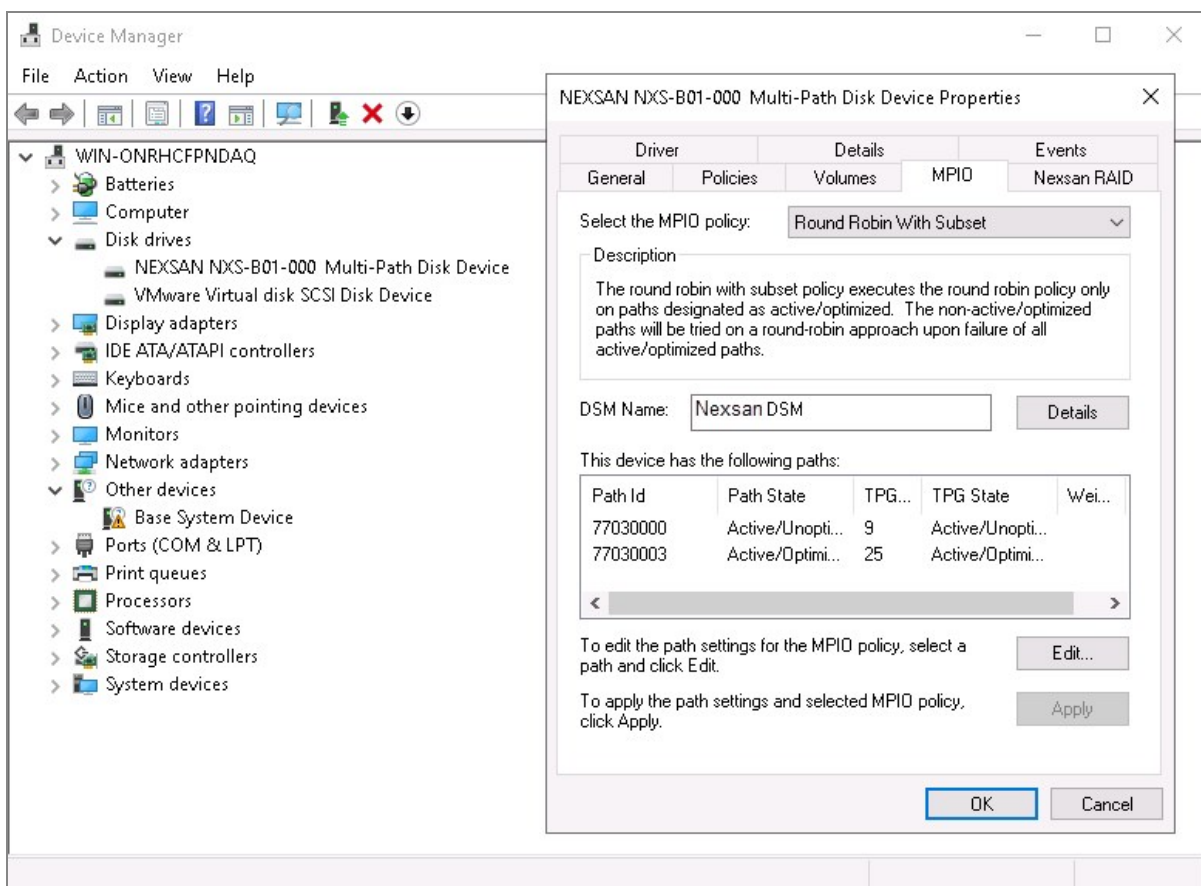
**Note** The following steps are applicable for Nexsan Storage Systems using either the Fibre Channel/SAS or iSCSI connection protocols.

▶ **To verify that storage paths are configured:**

- Open the Device Manager:
  - In Windows Server 2016 or later, click **Search > Device Manager**.
  - In Windows Server 2012 R2 or earlier, click **Settings > Device Manager**.
- Expand *Disk drives* and confirm that the Nexsan volumes are listed.

3. On each Nexsan disk, right-click and select **Properties**.

Figure 2-6: MPIO tab in Device Manager Multi-Path Disk Device Properties



4. In the *Disk Device Properties* dialog, click the **MPIO** tab to
  - a. Review the *Select the MPIO policy* field. Make a policy selection according to your requirements. See "Load balancing policies" ([page 12](#)).
  - b. Review the *DSM Name* field. On Windows Server 2016 or later, it should be *Microsoft DSM*. On Windows Server 2012 R2 or earlier, it should be *Nexsan DSM*.
5. Click **OK** to close the *Disk Device Properties* dialog.

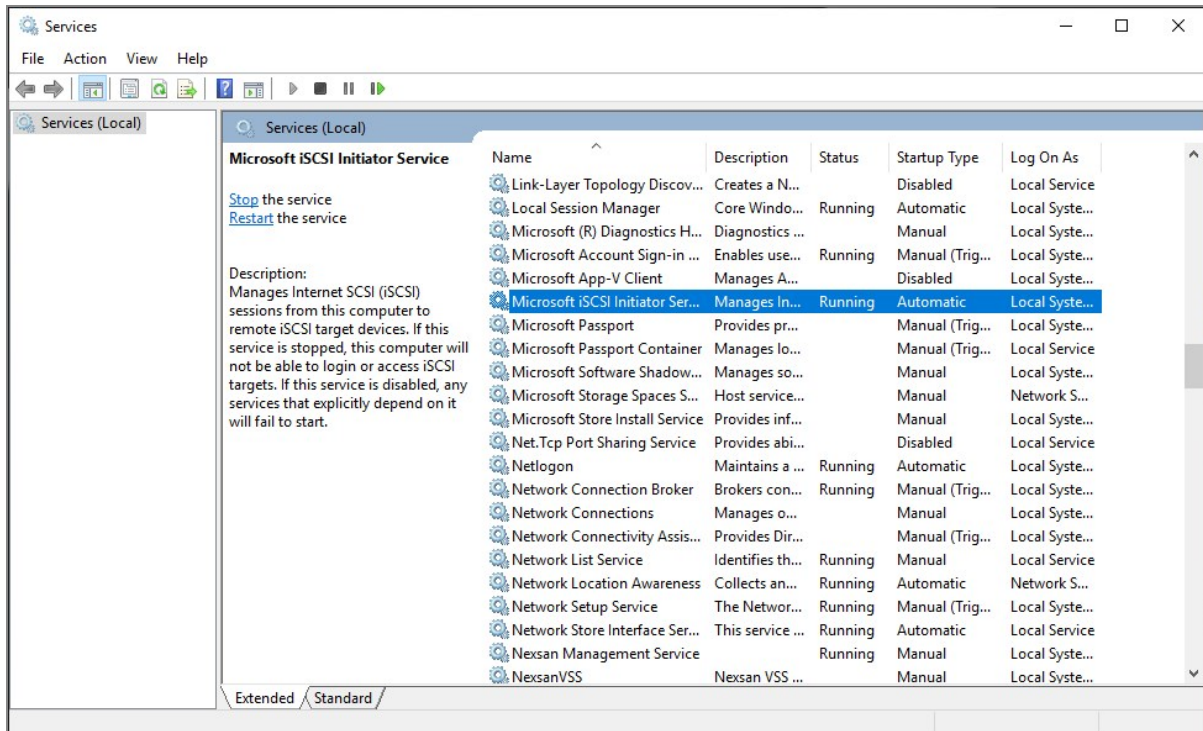
## Enabling the iSCSI Initiator service

For iSCSI networks, whether 1Gb/s or 10Gb/s, use the following additional steps if necessary to enable the iSCSI Initiator service.

► **To start Microsoft iSCSI Initiator service:**

1. Open the *Services* dialog:
  - a. In Windows 2016 and later, click **Start > Windows Administrative Tools > Services**, or
  - b. In Windows 2012 R2 or earlier, click **Start > Administrative Tools > Services**

Figure 2-7: Windows Server *Services* dialog



2. In the *Services* window, scroll down until **Microsoft iSCSI Initiator** is visible, then double-click it. The *Microsoft iSCSI Initiator Service Properties* dialog opens.
3. Under *Service Status*, click the **Start** button. The Microsoft iSCSI Initiator Service starts.
 

**Note** If the service is running, the button is greyed out.
4. Close the *Services* dialog.


## Connecting to a target storage device over iSCSI

For iSCSI connections, whether 1Gb/s or 10Gb/s, use the steps below to connect from Windows to a target Nexsan Storage System.

► **To connect to a target storage device:**

1. Open iSCSI Initiator:
  - a. In Windows Server 2016 or later, click **Start > Windows Administrative Tools > iSCSI Initiator**.  
or
  - b. In Windows Server 2012 R2 or earlier, click **Start > Administrative Tools > iSCSI Initiator**.
2. On the **Targets** tab, enter the IP address of an iSCSI port (not the management port) in the **Target** text field, then click **Quick Connect**.

Figure 2-8: *Discovered target* list in the *iSCSI Initiator Properties* dialog



3. On the *Quick Connect* dialog, verify that the storage device is in the *Discovered targets* list. The *iSCSI Initiator Properties* dialog displays the connected storage device in the *Discovered targets* list.



4. Click **OK** to close iSCSI Initiator. To verify that the multipathing is enabled, see "To verify that storage paths are configured:" ([page 19](#)).

## Recommended disk timeout settings

To allow sufficient time for recovery after a rolling restart, it is recommended that timeouts be set. For details, see the *Nexsan Nexsan High-Density Storage User Guide*.

*Windows Server 2012 R2 and earlier*

If the timeouts are set to a lower value than 150 seconds, the first time you run Nexsan MSIO, a warning dialog will display:

Figure 2-9: Nexsan MSIO disk timeout warning



To fix the problem, check the **Attempt to fix the above warnings** check box and click **OK**. The system reboots, and the timeout values are set to the appropriate settings.

*Windows Server 2016 or later*

In Windows Server 2016 or later, Nexsan Storage uses the default disk timeout value of 60 seconds.

### ► To view or update MPIO Settings in Windows Server 2016 or later:

1. Open Windows Server PowerShell, and run Get-MPIOSetting.

Example output:

```
PS C:\> Get-MPIOSetting
PathVerificationState : Disabled
PathVerificationPeriod : 30
PDORemovePeriod : 20
RetryCount : 3
RetryInterval : 1
UseCustomPathRecoveryTime : Disabled
CustomPathRecoveryTime : 40
DiskTimeoutValue : 60
```

2. You can update the timeout setting using Set-MPIOSetting.

Example:

```
Set-MPIOSetting -NewDiskTimeout 100
```





## Multipathing on Linux

---

This chapter contains the following sections:

Prerequisites .....	25
Multipathing components and definitions .....	25
Configuring Linux multipathing .....	28
Multipath.conf examples .....	31

### Prerequisites

In order to enable multipathing on Linux hosts, the following software packages must be installed:

- RHEL / Centos device-mapper-multipath (including sysvinit, libs and devel)
- Open SUSE – yast2-multipath, multipath-tools.
- Debian / Ubuntu – multipath-tools (including boot and initramfs)
- sg3\_utils (all distributions)
- lsscsi (all distributions)
- Up-to-date driver package for your interface card (Fibre Channel, iSCSI, and SAS ports)

### Multipathing components and definitions

Here are the three multipathing components and some brief definitions:

multipath	The multipath binary / tool is used to build the multipath map based on the LUNs attached to a Linux host and the multipath configuration.
multipath.conf	A text-based configuration file which is used to configure and build up the multipath map and multipathed devices.
multipathd	The multipath daemon monitors for path related events and takes action dependent upon the multipath map and configuration.

---

## multipath.conf

The `multipath.conf` file has the following sections, which are used to build up a multipath map using different attributes and values:

- defaults
- blacklist
- blacklist\_exceptions
- devices
- multipaths

Nexsan recommends creating a NEXSAN device instance for the multipath configuration.

Example device:

```
device {  
  
    vendor                "NEXSAN**"  
    product               "**"  
    path_grouping_policy  group_by_prio  
    path_checker          tur  
    path_selector         "round-robin 0"  
    prio                  alua  
    features              "1 queue_if_no_path" <May be deprecated in some  
                        distributions.  
    failback              immediate  
    rr_weight             priorities  
    no_path_retry         queue  
    fast_io_fail_tmo      15  
    dev_loss_tmo          150 (fibre transport only)  
    rr_min_io             1000  
    rr_min_io_rq          1  
  
}
```

The string "NEXSAN\*\*" is used for the vendor attribute because some Linux distributions require the vendor identifier to be eight characters long. Some multipath packages include a `multipath.conf` template.

Centos example:

```
/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf
```

A `man` page is available for further information about the `multipath.conf` contents. For more granularity or control, use the `multipath` section to define the configuration per multipath device. The WWID identifier will be needed for each multipath instance.

## Example

```

multipath {
    wwid 36000402005d801217274dd340000000
    0
    alias E18V-LUN1
    path_grouping_policy group_by_prio
    ....
}

```

## Attributes and Values

Attribute	Value	Description
vendor	"NEXSAN*"	Used to identify Nexsan systems.
product	"*"	Appended to the vendor attribute because some Linux distributions require the vendor identifier to be eight characters long.
prio	alua	Sets the path priority using the SCSI SPC-3 ALUA bit. A higher priority value is set for the optimized paths, which will be through the Controller Target ports which own the Array / LUN.
path_selector	"round-robin 0"	Sets the multipath policy to round robin.
rr_weight	priorities	The round-robin policy will set the path priority values / weights (as set by the prio alua).
path_grouping_policy	group_by_prio	Paths with the same priority are grouped together.
path_checker	tur	Sets the path checker to use the SCSI command Test Unit Ready (TUR).
failback	immediate	Ensures immediate failback to a max priority path.
fast_io_fail_tmo	15	Time taken in seconds before I/O to a path if failed if there is a fault is found with the fibre channel transport.
dev_loss_tmo	150	Time taken in seconds before a device is removed from the host if it cannot communicate with the device.
no_path_retry	queue	I/O is queued for a multipathed device if no paths are available.
features	"1 queue_if_no_path"	Works in the same way as <code>no_path_retry queue</code> but may be obsolete in some distributions.

See also:

[Configuring Linux multipathing](#) on the next page

[Multipath.conf examples on page 31](#)

## Configuring Linux multipathing

Use the steps in this section to configure multipathing for Linux with Nexsan Storage Systems.

► **To configure multipathing for Linux:**

1. Using a text editor, add the Nexsan configuration to the device section of the `multipath.conf` file:

```
device {
    vendor                "NEXSAN*"
    product               "*"
    path_grouping_policy  group_by_prio
    path_checker          tur
    path_selector         "round-robin 0"
    prio                  alua
    features              "1 queue_if_no_path" <May be deprecated in some
                        distributions.
    failback              immediate
    rr_weight             priorities
    no_path_retry         queue
    fast_io_fail_tmo     15
    dev_loss_tmo         150 (fibre transport only)
    rr_min_io             1000
    rr_min_io_rq         1
}
```

The device-specific instance above will apply to all attached Nexsan Storage Systems.

If a different configuration is needed for a SATABoy or SATABeast storage systems, create a new device instance and change the product identifier to "SATAB.\*".

For E-Series storage systems, use the product identifier "NXS-B0.\*".

2. Save the `multipath.conf` file and ensure that the multipath configuration includes the Nexsan device:

```
linux-cwni:/etc # multipath -t
...
device {
    vendor "NEXSAN*"
    product "*"
}
```

```
path_grouping_policy "group_by_prio"
path_selector "round-robin 0"
```

...

3. Run the `lsscsi` command to list which Nexsan devices are attached to the Linux host. In the example below, the host is connected to four Target ports of an E-18V system without any LUNs mapped:

```
linux-cwni:~ # lsscsi -g | grep NEXSAN
[3:0:0:0] disk NEXSAN NXS-B01-000 Q011 -
[4:0:0:0] disk NEXSAN NXS-B01-000 Q011 -
[5:0:0:0] disk NEXSAN NXS-B01-000 Q011 -
[6:0:0:0] disk NEXSAN NXS-B01-000 Q011 -
```

Each of the Nexsan disk devices above represents one of the Target ports. Note the difference in the `[host:bus:target:lun]` address for each Target port.

If LUNs have been mapped to the Target ports and are expected in the list, run the rescan script included with the `sg3_util` package.

Here is the same E-18V system with LUN 1 mapped to the four Target ports:

```
linux-cwni:~ # lsscsi -g | grep NEXSAN
[3:0:0:0] disk NEXSAN NXS-B01-000 Q011 -
[3:0:0:1] disk NEXSAN NXS-B01-000 Q011 /dev/sdb
[4:0:0:0] disk NEXSAN NXS-B01-000 Q011 -
[4:0:0:1] disk NEXSAN NXS-B01-000 Q011 /dev/sdc
[5:0:0:0] disk NEXSAN NXS-B01-000 Q011 -
[5:0:0:1] disk NEXSAN NXS-B01-000 Q011 /dev/sdd
[6:0:0:0] disk NEXSAN NXS-B01-000 Q011 -
[6:0:0:1] disk NEXSAN NXS-B01-000 Q011 /dev/sde
```

There are now four SCSI Drives representing each of the mapped LUNs.

If more LUNs are mapped to the Target ports, additional SCSI Drives will be listed under each Target port device (the LUN number being the end digit). For example:

```
[6:0:0:2] disk NEXSAN NXS-B01-000 Q011 /dev/sdf
```

You may wish to review the configuration using the GUI of the attached Nexsan Storage System to verify that the expected Nexsan devices are visible to the system.

4. Ensure that the `dm_multipath` module is loaded using `lsmod | grep dm_multipath`

If it is not loaded, run the `modprobe` command to load the module:

```
modprobe dm_multipath
```

5. Clear the existing map using the `multipath -F` command.

**Note** The `-F` flag clears all multipathing mapping, which may disrupt existing multipathed devices. If you have existing multipathed devices, use `multipath -f` to clear the configuration for the new volumes only.

6. Build a new map using the `multipath -v2` command, as follows:

```
linux-cwni:/etc # multipath -v2
```

```

create: 36000402005d801217274dd3400000000 undef NEXSAN,NXS-B01-000
size=3.6T features='0' hwhandler='1 alua' wp=undef
|-+ policy='service-time 0' prio=50 status=undef
||- 3:0:0:1 sdb 8:16 undef ready running
|`- 4:0:0:1 sdc 8:32 undef ready running
`-+ policy='service-time 0' prio=10 status=undef
|- 5:0:0:1 sdd 8:48 undef ready running
`- 6:0:0:1 sde 8:64 undef ready running

```

7. Verify the new multipath map using the `multipath -ll` command:

```

linux-cwni:/etc # multipath -ll
36000402005d801217274dd3400000000 dm-0 NEXSAN,NXS-B01-000
size=3.6T features='0' hwhandler='1 alua' wp=rw
|-+ policy='service-time 0' prio=50 status=active
||- 3:0:0:1 sdb 8:16 active ready running
|`- 4:0:0:1 sdc 8:32 active ready running
`-+ policy='service-time 0' prio=10 status=enabled
|- 5:0:0:1 sdd 8:48 active ready running
`- 6:0:0:1 sde 8:64 active ready running

```

Ensure that for each multipathed device (dm-0 in this example):

- a. A WWID is assigned and the Vendor / Product strings are listed:

```
36000402005d801217274dd3400000000 dm-0 NEXSAN,NXS-B01-000
```

- b. The number of paths and SCSI Drives are correct.
- c. The paths are grouped by a higher (`prio=50`) and lower (`prio=10`) values. The higher prio value is used for the optimized paths (paths through the owning Controller)
- d. The status is enabled and each path is active ready.

**Note** If the map is incorrect, try recreating it using the `multipath -v3` command. Look for errors in the output or incorrect attribute values.

8. Start the `multipathd` daemon:

```

systemctl start multipathd (SystemD)
/etc/init.d/multipathd start (SysV)

```

9. Monitor any path events using the messages file.

In the example below, the device with the `major:minor` number (sdb) has failed for dm-0 and is later reinstated:

```

[247391.992993] device-mapper: multipath: Failing path 8:16.
[247551.102072] sd 3:0:0:1: Warning! Received an indication that the LUN assignments on this target
have changed.
[247552.103354] device-mapper: multipath: Reinstating path 8:16.
[247552.113589] sd 3:0:0:1: alua: port group 05 state A preferred supports toluSNA

```

```
[247552.114257] sd 3:0:0:1: alua: port group 05 state A preferred supports toluNA
```

## Multipath.conf examples

**Note** Changes were made to the format of the path identifier (`getuid`) and priority (`prio_callout`) attributes between multipath modules 0.4.8 and 0.4.9. While these examples are as accurate as possible, the exact settings required may vary between Linux distributions and versions.

### Redhat 5/Centos

```
defaults {
    udev_dir                /dev
    polling_interval        10
    selector                "round-robin 0"
    path_grouping_policy    group_by_prio
    getuid_callout          "/sbin/scsi_id -g -u -s /block/%n"
    prio_callout            "/sbin/mpath_prio_alua %d"
    path_checker            directio
    rr_min_io               1000
    rr_weight               priorities
    failback                immediate
    no_path_retry           15
    user_friendly_names     yes
}

blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z] [[0-9]*]"
    devnode "^cciss!c[0-9]d[0-9]*"
}

devices {
    device {
        vendor                "NEXSAN  "
        product               "NXS-B0.*"
        path_grouping_policy  group_by_prio
        getuid_callout        "/sbin/scsi_id -g -u -s /block/%n"
        prio_callout          "/sbin/mpath_prio_alua %d"
        path_checker          directio
        path_selector         "round-robin 0"
        hardware_handler      "0"
        rr_weight              priorities
        no_path_retry         15
        failback               immediate
    }
}

multipaths {
    multipath {
        wwid                  36000402000d8129f66a43a5e0000000
        alias                 E60
        path_grouping_policy  group_by_prio
    }
}
```

```

        path_selector          "round-robin 0"
        prio_callout          "/sbin/mpath_prio_alua %d"
        failback              immediate
        rr_weight             priorities
        no_path_retry         15
    }
}

```

## Redhat 6

```

defaults {
    udev_dir                /dev
    polling_interval        10
    selector                "round-robin 0"
    path_grouping_policy    group_by_prio
    getuid_callout          "/lib/udev/scsi_id --whitelisted"
--device=/dev/%n"
    prio                    alua
    path_checker            directio
    rr_min_io               1000
    rr_weight               priorities
    failback               immediate
    no_path_retry           15
    user_friendly_names     yes
}

blacklist {
    devnode "^ (ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z] [[0-9]*]"
    devnode "^cciss!c[0-9]d[0-9]*"
}

devices {
    device {
        vendor                "NEXSAN "
        product               "NXS-B0.*"
        path_grouping_policy  group_by_prio
        getuid_callout        "/lib/udev/scsi_id --whitelisted"
--device=/dev/%n"
        prio                  alua
        path_checker          directio
        path_selector         "round-robin 0"
        hardware_handler      "0"
        rr_weight             priorities
        no_path_retry         15
        failback              immediate
    }
}

multipaths {
    multipath {

```



```

        wwid                36000402000d8129f66a43a5e00000000
        alias                E60
        path_grouping_policy group_by_prio
        path_selector        "round-robin 0"
        prio                 alua
        failback             immediate
        rr_weight            priorities
        no_path_retry        15
    }
}

```

## DEBIAN Lenny

```

defaults {
    udev_dir                /dev
    polling_interval        10
    selector                "round-robin 0"
    path_grouping_policy    group_by_prio
    getuid_callout         "/lib/udev/scsi_id --whitelisted -g -u /dev/%n"
    prio                   mpath_prio_alua /dev/%n
    path_checker            directio
    rr_min_io              1000
    flush_on_last_del      no
    max_fds                 8192
    rr_weight              priorities
    failback               immediate
    no_path_retry          15
    user_friendly_names    yes
    mode                   644
    uid                    0
    gid                    disk
}

blacklist {
    wwid 26353900f02796769
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^(hd[a-z]*)((0-9)*)"
    devnode "sda"
}

devices {
    device {
        vendor                "NEXSAN  "
        product               "NXS-B0.*"
        path_grouping_policy  group_by_prio
        getuid_callout        "/lib/udev/scsi_id -g -u --whitelisted /dev/%n"
        path_checker          directio
        path_selector         "round-robin 0"
        hardware_handler      "0"
        failback              immediate
        rr_weight             priorities
    }
}

```

```

        no_path_retry          15
        product_blacklist     LUNZ
    }
}

multipaths {
    multipath {
        wwid                   36000402000d8129f66a43a5e00000000
        alias                   E60
    }
}

```

## DEBIAN Wheezy

```

defaults {
    udev_dir                   /dev
    polling_interval           10
    path_selector               "round-robin 0"
    # getuid_callout            "/lib/udev/scsi_id --whitelisted --device=/dev/%n"
    path_grouping_policy       multibus
    prio                        const
    path_checker                tur
    rr_min_io                   1000
    rr_weight                   priorities
    failback                    immediate
    no_path_retry               5
    flush_on_last_del           no
    user_friendly_names         yes
    bindings_file               /etc/multipath/bindings
}

blacklist_exceptions {
    device {
        vendor XYRATEX*
    }

    device {
        vendor transtec*
    }

    device {
        vendor IFT*
    }

    device {
        vendor LSI*
    }

    device {
        vendor NEXSAN
    }
}

```

```

devices {
    device {
        vendor          "XYRATEX"
        product         "F5412E"
        path_grouping_policy failover
    }

    device {
        vendor          "transtec"
        product         "PROVIGO 550F"
        path_grouping_policy failover
    }

    device {
        vendor          "IFT"
        product         "S24F-R1840-4"
        path_grouping_policy group_by_prio
        hardware_handler "1 alua"
        prio            alua
    }

    device {
        vendor          "IFT"
        product         "S16F-R1430"
        path_grouping_policy group_by_prio
#       hardware_handler "1 alua"
        hardware_handler "0"
        prio            alua
        rr_weight       uniform
        no_path_retry   12
    }

    device {
        vendor          "LSI"
        product         "INF-01-00"
    }

    device {
        vendor          "NEXSAN "
        product         "NXS-B0.*"
        path_checker    directio
        prio            alua
        path_grouping_policy group_by_prio
        hardware_handler "1 alua"
        no_path_retry   5
#       no_path_retry   2
#       fast_io_fail_tmo 6
    }
}

```

## SUSE 11 (including OpenSUSE)

```

defaults {
    udev_dir                /dev
    polling_interval        10
    path_selector           "round-robin 0"
    path_grouping_policy    group_by_prio
    getuid_callout          "/lib/udev/scsi_id --whitelisted -g -u -d /dev/%n"
    prio                    alua
    path_checker            directio
    rr_min_io               1000
    flush_on_last_del       no
    max_fds                 8192
    rr_weight               priorities
    failback                immediate
    no_path_retry           15
    user_friendly_names     yes
    mode                    644
    uid                     0
    gid                     disk
}

blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st) [0-9]*"
    devnode "^hd[a-z] [[0-9]*]"
}

devices {
    device {
        vendor                "NEXSAN  "
        product               "NXS-B0.*"
        path_grouping_policy    group_by_prio
        getuid_callout          "/lib/udev/scsi_id --whitelisted -g -u - d /dev/%n"
        prio                    alua
        path_checker            directio
        path_selector           "round-robin 0"
        hardware_handler        "0"
        failback                immediate
        rr_weight               priorities
        no_path_retry           15
        product_blacklist       LUNZ
    }
}

multipaths {
    multipath {
        wwid                   36000402000d8129f66a43a5e0000000
        alias                   E60
        path_selector           "round-robin 0"
        failback                immediate
        rr_weight               priorities
    }
}

```

```
        no_path_retry          15
    }
}
```



## Multipathing on Solaris

---

The Solaris operating system multipathing feature, MPxIO, was formerly known as the StorageTek Traffic Manager Software (STMS). MPxIO provides the functionality to manage and organize multiple paths to the same storage volume. MPxIO supports SAS, Fibre Channel and iSCSI transports (support may vary between distributions).

By default, MPxIO will work with StorageTek storage units. Extra configuration is required to support Nexsan units. For the default `scsi_vhci.conf` file (`/kernel/drv/scsi_vhci.conf`) to work with Nexsan units, the proper Vendor and Product settings need to be added to the file's `device-type-scsi-options-list`.

This chapter contains the following sections:

<a href="#">Configuring Solaris I/O multipathing</a> .....	41
<a href="#">Rolling restarts and configurable timeouts</a> .....	45

### Prerequisites

- Solaris 10, Update 3 or later must be installed on the host.
- The `scsi_vhci` module must be loaded, that is:

```
modinfo -c | grep scsi_vhci:  
  
31      1 scsi_vhci      LOADED/INSTALLED
```

- Root user access must be available.
- The LUN mapping and host-access must be set up on the Nexsan system.
- The host must be logged into the Nexsan system targets.

---

## Multipath components

Here is a list of multipathing components and descriptions:

Component	Description
<code>scsi_vhci.conf</code>	Configuration file for the <code>scsi_vhci</code> driver, which supports multipathed devices.
<code>mpathadm</code>	Command to discover and administer multipathed devices.
<code>stmsboot</code>	Command to switch multipath on or off globally, or by controller type, that is iSCSI, fibre channel, or SAS.
<code>sd.conf</code>	Configuration file for the SCSI device driver which is used to set the <code>sd</code> timeout, and SCSI device tunables.
<code>iscsi.conf</code>	Configuration file for the iSCSI driver.
<code>fc.conf</code>	Configuration file for the fibre channel driver.

## Multipathed device GUID

Solaris I/O multipathing (previously known as MPXIO) is enabled by default and uses the `scsi_vhci` driver to manage multipathed devices.

Each multipathed device is identified with a (raw) GUID in the following format:

```
/dev/rdisk/c0t6000402003E0175B77BCFB0F00000000d0
```

The GUID is referenced as the STMS device name and the paths for each device are referenced as the non-STMS device name.



## Configuring Solaris I/O multipathing

Use this procedure to configure I/O multipathing on Solaris.

► **To configure I/O multipathing on Solaris:**

1. Edit the `/kernel/drv/scsi_vhci.conf` file and add the Nexsan Vendor and Product strings.

Example:

```
#
device-type-scsi-options-list =
"NEXSAN NXS-B01-000", "enable-option";
#

enable-option = 0x1000000;

#
```

Note, the NEXSAN Vendor ID (VID) has two spaces added to make the string 8 characters in length.

2. Check that `MPXIO` is enabled in the transport driver configuration file.

Example:

```
/kernel/drv/fp.conf (fibre channel driver)

mpxio-disable="no";
```

**Note** If `mpxio-disable` is set to yes, edit the driver configuration file and set the property to no.

A reboot of the host will be required.

3. Run the `mpathadm` command to list the multipathed devices and associated paths.

Example of two multipathed devices each with two available paths:

```
# mpathadm list LU

/dev/rdisk/c0t6000402003E0175B77BCFB0F00000000d0s2
    Total Path Count: 2
    Operational Path Count: 2

/dev/rdisk/c0t6000402003E0175B77AE818500000000d0s2
    Total Path Count: 2
    Operational Path Count: 2
```

4. To show more detail for each multipathed device and paths, run the `mpathadm show LUN` command.

Example for a multipathed device connected via fibre channel transport:

```
root@openindianaFC:~# mpathadm show LUN
/dev/rdisk/c0t6000402003E0175B77BCFB0F00000000d0s2

Logical Unit:  /dev/rdisk/c0t6000402003E0175B77BCFB0F00000000d0s2

mpath-support:  libmpscsi_vhci.so
```

Vendor: NEXSAN  
Product: NXS-B01-000  
Revision: R011  
Name Type: unknown type  
Name: 6000402003e0175b77bcfb0f00000000  
Asymmetric: yes  
Current Load Balance: round-robin  
Logical Unit Group ID: NA  
Auto Failback: on  
Auto Probing: NA  
Paths:  
    Initiator Port Name: 10000000c93e4bc0  
    Target Port Name: 5000402103e7175b  
    Override Path: NA  
    Path State: OK  
    Disabled: no  
  
    Initiator Port Name: 10000000c93e4bc0  
    Target Port Name: 5000402003e7175b  
    Override Path: NA  
    Path State: OK  
    Disabled: no

```

Target Port Groups:
  ID: 2
  Explicit Failover: no
  Access State: active optimized
  Target Ports:
    Name: 5000402103e7175b
    Relative ID: 2
  ID: 1
  Explicit Failover: no
  Access State: active optimized
  Target Ports:
    Name: 5000402003e7175b
    Relative ID: 1

```

**Notes:**

- The Vendor and Product IDs are listed as NEXSAN NXS-B01-000 which matches the `scsi_vhci.conf` driver file.
  - The Asymmetric attribute is set to Yes and the two paths have an active optimized Access State, due to the paths being connected to the owning controller of the Nexsan system. An active non-optimized Access State is reported for paths to the non-owning controller.
  - The load balancing policy is set to round-robin.
5. Run the `stmsboot` command to enable multipathing, either globally or for the specific controller type.

Example to enable multipathing globally:

```

:~# stmsboot -e

```

```

WARNING: stmsboot operates on each supported multipath-capable controller
detected in a host. In your system, these controllers are

```

```

/pci@0,0/pci8086,244e@1e/pci10df,fa00@1,1/fp@0,0

```

```

/pci@0,0/pci8086,244e@1e/pci10df,fa00@1/fp@0,0

```

```

If you do NOT wish to operate on these controllers, please quit stmsboot
and re-invoke with -D { fp | mpt | mpt_sas | pmcs} to specify which
controllers you wish to modify your multipathing configuration for.

```

6. Reboot the host when prompted.
7. Run the `stmsboot` command to list the multipathed device and associated paths.  
Example showing a multipathed device (STMS device name) and the SCSI devices / paths (non-STMS device name)

```
:~# stmsboot -L
```

```
non-STMS device name          STMS device name
```

```
-----
```

```
/dev/rdisk/c5t5000402103E7175Bd21
```

```
/dev/rdisk/c0t6000402003E0175B77BCFB0F00000000d0
```

```
/dev/rdisk/c5t5000402003E7175Bd21
```

```
/dev/rdisk/c0t6000402003E0175B77BCFB0F00000000d0
```

### Notes:

Devices connected using iSCSI transport will have a different format from those shown above. For example,

```
/dev/rdisk//iscsi/disk@0000iqn.1999-02.com.nexsan%3Ap1%3Anxs-b01-
```

```
000%3A03e5175b0002,22 /dev/rdisk/c0t6000402003E0175B77AE818500000000d0
```

```
/dev/rdisk/c7t5d0
```

```
/dev/rdisk/c0t6000402003E0175B77AE818500000000d0
```

If `stmsboot - L` shows no devices, it may be necessary to switch off MPIOX in the relevant driver file, reboot, re-enable MPIOX and reboot again.

## Rolling restarts and configurable timeouts

To support a rolling restart of the controllers, it is recommended to set the host-side timeout to 150 seconds. This will ensure no loss of service while the controllers are rebooted and the targets are offline. The restart usually takes the storage targets offline for between 30 to 60 seconds.

### SCSI device timeout and retry

Each SCSI Device (sd) has associated command timeout and retry values, which are by default set to 60 seconds and 3 retries, respectively.

The command timeout and retry values are set in the `sd.conf` driver file using the `sd_io_time` and `sd-config-list` properties:

```
sd_io_time - command timeout
sd-config-list = "NEXSAN NXS-B01-000", "retries-timeout = x
```

The `sd-config-list` property uses the Vendor / Product IDs to match the Nexsan system.

For a Rolling Restart, the recommended values are:

```
sd_io_time = 150;
sd-config-list = "NEXSAN NXS-B01-000", "retries-timeout = 5";
```

#### Notes:

- The `sd` timeout values may need to be tailored for different environments.
- For SPARC fibre attached systems, the timeout and retry values are set in the `ssd.conf` driver file.

### Transport timeouts

There are configurable fibre and iSCSI properties, which are discussed below.

#### Fibre Channel transport

The `fp_offline_ticker` and `fcp_offline_delay` properties determine how long to wait until a fibre link reports error recovery or a SCSI device/path is taken offline.

The properties are set in the following files:

```
/kernel/drv/fp.conf (fp_offline_ticker=xxx;) (default = 90s)
/kernel/drv/fcp.conf (fcp_offline_delay=xxx;)
```

For a Rolling Restart, it is recommended to set the properties to

```
fp_offline_ticker=150;
fcp_offline_delay=60;
```

## iSCSI transport

There are three initiator or target tunable properties. They determine:

1. The time to wait for a response to a log-in command.
2. The time to keep re-trying to establish a failed connection.
3. The time between retries.

The properties are listed or modified using the `iscsiadm` command.

Example:

```
bash-3.00# iscsiadm list initiator-node

Initiator node name: iqn.1986-02 ...

...

Tunable Parameters (Default/Configured): Session Login Response
Time: 60/- Maximum Connection Retry Time: 180/- Login Retry Time
Interval: 60/-

bash-3.00# iscsiadm modify initiator-node -T polling-login-delay=150
```

**Note** The following names (shown in italics) must be used with the `-T` parameter to modify the tunable values:

- Session Login Response Time: *recv-login-rsp-timeout*
- Maximum Connection Retry Time: *conn-login-max*
- Maximum Connection Retry Time: *polling-login-delay*

No changes to the default values are recommended.

```
fp_offline_ticker=150;
fcp_offline_delay=60;
ssd_io_time=150; (may not be required for AMD / x86)
ssd_retry_count=5;
sd_io_time=150;
sd_retry_count=5;
```

## Multipathing on IBM AIX

---

This chapter contains the following sections:

IBM AIX Multiple Path IO .....	47
Nexsan Multi-Path I/O for IBM AIX .....	49
Installing MSIO software for IBM AIX .....	49
Discover MSIO volumes .....	50
Viewing and changing MSIO policy .....	50
Volume groups .....	52
HACMP best practices .....	52
Uninstalling MSIO .....	52

### IBM AIX Multiple Path IO

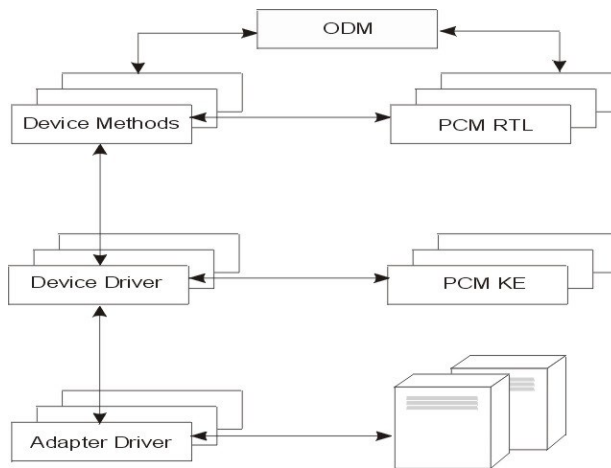
With IBM AIX Multiple Path I/O (MPIO), a device can be uniquely detected through one or more physical connections, or paths.

A path-control module (PCM) provides the path management functions.

An MPIO-capable device driver can control more than one type of target device. A PCM can support one or more specific devices. Therefore, one device driver, such as the Nexsan MSIO, can be interfaced to multiple PCMs that control the I/O across the paths to each of the target devices.

This illustration shows the interaction between the different components that make up the MPIO solution. In this figure, the MPIO device driver controls multiple types of target devices, each requiring a different PCM (KE=Kernel Extension, RTL=Run-Time Loadable).

Figure 5-1: MPIO component interaction



The AIX PCM consists of the PCM RTL configuration module and the PCM KE kernel extension. The PCM KE supplies path-control management capabilities to any device driver that supports the MPIO interface.

The PCM KE depends on device configuration to detect paths and communicate that information to the device driver. The PCM KE can provide more than one routing algorithm, which can be selected by the user. The PCM KE also helps collect information that can be used to determine and select the best path for any I/O request. The PCM KE can select the best path based on a variety of criteria, including load balancing, connection speed, connection failure, etc.

The AIX PCM has a health-check capability that can be used to do the following:

- Check the paths and determine which paths are currently usable for sending I/O.
- Enable a path that was previously marked as having failed because of a temporary path fault (for example, when a cable to a device was removed and then reconnected).
- Check currently unused paths that would be used if a failover occurred (for example, when the algorithm attribute value is `failover`, the health check can test the alternate paths).



## Nexsan Multi-Path I/O for IBM AIX

Nexsan MultiSan IO (MSIO) support for AIX is based on the AIX standard Multi-Path driver. The only thing required for MSIO support is a definition of appropriate objects in the AIX ODM database.

Nexsan MSIO provides the following load balancing policies:

- **Failover** — No load balancing is performed. The application specifies a primary path and a set of standby paths. The primary path is used for processing device requests. If the primary path fails, one of the standby paths is used. Standby paths must be listed in decreasing order of preference (the most preferred path first).
- **Round-robin** — Uses all available paths for I/O in a balanced, round-robin fashion.

## Installing MSIO software for IBM AIX

It is recommended to keep the entire installation set in a separate directory so that the installation process can start with the creation of this directory and the copying of four files:

```
extract perform:
# mkdir /opt/MSIO
# cd /opt/MSIO
```

Copy the supplied NEXSAN-MSIO\_AIX6.1 file into the MSIO folder. The MSIO folder should contain the following files:

- install.sh
- disk.fcp.nexsan.add
- nexsan.msg
- pcm.friend.nexsan.add
- uninstall.sh

To install Multi-Path I/O support on the AIX machine, execute the install.sh script from the MSIO folder:

```
# ./install.sh
```

This script will execute the following commands from the MSIO folder path:

```
# gencat nexsan.cat nexsan.msg
# cp nexsan.cat /usr/lib/methods
# odmadd disk.fcp.nexsan.add
# odmadd pcm.friend.nexsan.add
# bosboot -a
```

The installation process is completed.

**Note** From now on, multipath support will be provided for each MSIO volume. The default policy for the entire class of multipath volumes in AIX is Failover. To change the Multi-path policy, please refer to [Viewing and changing MSIO policy](#) on the next page.

## Discover MSIO volumes

Before you start using Nexsan MSIO volumes on an AIX machine, execute the following:

```
# lsdev -c disk
```

You may see Nexsan MSIO volumes recognized by the AIX machine as Other FC SCSI Disk Drive. To verify whether these are Nexsan volumes, execute:

```
# lsattr -El hdisk4 | grep lun_id
```

and match the printed LUN IDs with those defined by the Nexsan device.

In this case, remove the device:

```
# rmdev -d -l <hdisk4>
```

and rescan the FC channels using the `cfgmgr` command or rebooting the AIX host:

```
# cfgmgr -l fcs0
```

```
# cfgmgr -l fcs1
```

Check your devices again using the `lsdev` command. You should see Nexsan volumes as NEXSAN FC SCSI Disk Drive, like this:

```
bash-3.00# lsdev -c disk
hdisk4 Available 07-09-19 NEXSAN FC SCSI Disk Drive
hdisk5 Available 07-09-19 NEXSAN FC SCSI Disk Drive
hdisk6 Available 07-09-19 NEXSAN FC SCSI Disk Drive
```

## Viewing and changing MSIO policy

To check what the current load balance policy is for Nexsan volumes, use the `lsattr` command, specifying the `-El` attributes. For example, for device `hdisk7`, enter the command:

```
# lsattr -El hdisk7|grep algorithm
algorithm fail_over PCM algorithm True
```

The example below shows the current MSIO policy for disk 10 that is set to the Round-Robin policy (recommended for HACMP implementations):

```
bash-3.00# lsattr -El hdisk10
PCM PCM/friend/nexsanpcm Standard AIX PCM False
PR_key_value f1f2f3f4f5f6f7f8 Private reservation key True
algorithm round_robin PCM algorithm True
clr_q no Queue error clearing True
dist_err_pcnt 0 PCM error limit for failed path,percent True
dist_tw_width 50 PCM time-window, msec True
hcheck_cmd test_unit_rdy PCM health-check command True
```

```

hcheck_interval 30 PCM health-check interval True
hcheck_mode nonactive PCM health-check mode True
location Disk location, user defined True
lun_id 0x20000000000000 LUN ID False
lun_reset_spt yes N/A True
max_transfer 0x100000 Maximal transfer block size True
node_name 0x2001000402fc5509 Node name False
pvid 00057db22ae33f8600000000000000000 Physical volume ID False
q_err no Queue error handling True
q_type simple Queue type True
queue_depth 255 Queue depth True
reassign_to 120 Reassign timeout, sec True
reserve_policy no_reserve Reservation policy True
rw_timeout 90 Read/Write timeout, sec True
scsi_id 0x370300 SCSI ID False
sn_location 081083 VPD location for dynamic tracking False
start_timeout 30 Start timeout, sec True
ww_name 0x5000402301fc5509 WW Name False

```

To change the load balance policy for Nexsan volumes, use the `chdev` command, specifying the `PR_key_value`, `reserve_policy`, and `algorithm` attributes.

For example, for device `hdisk7`, enter the command:

```
# chdev -l hdisk7 -aPR_key_value=0xf1f2f3f4f5f6f7f8 -areserve_
policy=PR_exclusive -a algorithm=round_robin
```

For more information about load balance policies, see [Appendix A, Load Balancing Policies](#) on page 1.

## Volume groups

Once the load balance policy is in place, you can create an AIX volume group that will make use of the multipathed devices.

Use the `mkvg` command to create the volume group, as shown in the following example:

```
bash-3.00# mkvg -s 512 -y <my-vol-group> <hdisk10>
```

Once created, take the group online using the `varyonvg` command:

```
bash-3.00# varyonvg <my-vol-group>
```

Verify volume group creation using the `lsvg` command:

```
bash-3.00# lsvg <my-vol-group>
```

## HACMP best practices

We recommend the following High Availability Cluster Multiprocessing best practices with Nexsan MSIO:

- **MPIO Reservation Policy:** When deploying the Nexsan MSIO in a HACMP cluster environment, define the `no_reserve` reservation policy.
- **MPIO Algorithm:** When deploying the Nexsan MSIO in a HACMP cluster environment, define the `round_robin` algorithm.
- **MPIO Health Check:** When deploying the Nexsan MSIO in a HACMP cluster environment, define the health check `enabled`.

Example:

```
# chdev -l hdisk7 -a hcheck_mode=enabled -a reserve_policy=no_
reserve -a algorithm=round_robin
```

## Uninstalling MSIO

► **To remove Nexsan Multi-Path I/O support from an AIX machine:**

- Execute the `uninstall.sh` script from the MSIO folder:

```
# ./uninstall.sh
```

The uninstall script performs the following operations:

```
rm /usr/lib/methods/nexsan.cat
odmdelete -q uniquetype=disk/fcp/nexsan -o PdDv
odmdelete -q uniquetype=disk/fcp/nexsan -o PdAt
odmdelete -q uniquetype=disk/fcp/nexsan -o PdAtXtd
odmdelete -q uniquetype=disk/fcp/nexsan -o PdPathAt
odmdelete -q uniquetype=PCM/friend/nexsanpcm -o PdDv
odmdelete -q uniquetype=PCM/friend/nexsanpcm -o PdAt
odmdelete -q uniquetype=PCM/friend/nexsanpcm -o PdAtXtd
```

```
odmdelete -q uniquetype=PCM/friend/nexsanpcm -o PdPathAt
```

You may also wish to remove the directory containing the installation set.



## Multipathing on XenServer

This chapter contains the following sections:

Configuring multipathing on XenServer .....	55
Multipathing.conf example .....	55

### Configuring multipathing on XenServer

Set up multipathing on XenServer according to the instructions set forth in the Citrix support article: “Multipathing Overview for XenServer”, located at <http://support.citrix.com/article/CTX118791>.

### Multipathing.conf example

**Note** This example is specifically for iSCSI and should not be used for fibre channel.

```
device {
    vendor          "NEXSAN  "
    product         "SATAB.*"
    path_grouping_policy group_by_prio
    getuid_callout  "/sbin/scsi_id -g -u -s /block-
/%n"
    features        "0"
    path_checker    tur
# path_checker     directio
    hardware_handler "0"
    failback        150
    rr_weight       uniform
    no_path_retry   5
    rr_min_io       250
}
device {
    vendor          "NEXSAN  "
    product         "NXS.*"
    path_grouping_policy group_by_prio
```

---

```
getuid_callout          "/sbin/scsi_id -g -u -s /block-
/%n"
features                "0"
path_checker            tur
# path_checker           directio
hardware_handler        "0"
failback                150
rr_weight               uniform
no_path_retry           5
rr_min_io               250
}
}
```



# Glossary

---

## 2

### 2-Port Active-Active (2PAA) mode

A system mode for Nexsan Storage Systems. In this mode, each RAID controller operates as an independent node, but only one port is active on each controller. The second port operates in passive mode. Port 0 is active on controller 0, and port 1 is active on controller 1. The volumes are mapped to the active port on their owning controller. When one controller fails, the passive port on the other controller activates and takes over the host port functions of the failed controller. In a switched environment, failover is completely transparent to the hosts.

## 4

### 4-Port Active-Active (4PAA) mode

A system mode for Nexsan Storage Systems. In this mode, each RAID controller operates as an independent node, and all ports are active. Port 0 is the primary port on controller 0, and port 1 is the primary port on controller 1. The volumes must be mapped to at least one port on its owning controller and to the secondary port on the other controller, which requires the host to be configured for multipathing. When one controller fails, the secondary port on the other controller takes on the host address of the primary port on the failed controller,

allowing host I/O to continue (see failover); the host sees the storage become active through its second path.

## A

### Active Directory

Microsoft's directory service, used by all supported Windows operating systems. Abbreviated "AD".

### All Ports All LUNs (APAL) mode

A system mode for Nexsan Storage Systems. In this mode, the entire system operates as a single node. The volumes can be mapped to any or all ports on both RAID controllers. When a controller fails, the ports on that controller become inaccessible. However, if the volumes are mapped to ports on the other controller as well (which requires the host to be configured for multipathing), they remain accessible to the host, which sees the storage become active through its second path.

### array

A linked group of one or more physical, independent hard disk drives. See also "RAID".

**B**

## bit

The smallest unit of digital data, representing a 0 or a 1. Abbreviated “b”.

## boot drive

The device from which a computer’s operating system is loaded. Typically, an internal hard disk drive (or one of several partitions on such a drive) is used for this purpose, but any attached storage device—such as an optical disc drive, a USB flash drive, or other attached storage—can be used.

## byte

A unit of data that is 8 bits long. Often used for alphanumeric characters. Abbreviated “B”.

**C**

## CHAP

Challenge-Handshake Authentication Protocol. A method of authenticating any two devices that wish to communicate securely over a network. CHAP uses a “shared secret”, which is a plain text key known by both devices, to authenticate the connection.

**D**

## DHCP

Dynamic Host Configuration Protocol. A communication protocol that centrally manages and automates the assignment of IP addresses on a network.

## DNS

See “Domain Name System”.

## Domain Name System

A program or computer server that implements a name-service protocol. It maps a human-recognizable identifier to a system-internal, often numeric, identification or addressing component (usually an IP address).

## Dual Controller Non-Redundant (DCNR) mode

A system mode for Nexsan Storage Systems. In this mode, both RAID controller are active, but each controller operates as an independent node, and all ports are independent from each other. The volumes can only be mapped to ports on the controller that owns the RAID. They become inaccessible if the controller fails.

## duplex

A communication system where data flows in both directions between two devices. There are two configurations. “Half duplex” provides communication in both directions, but not at the same time; when one device transmits, the other device can only receive, and vice versa. For example, walkie-talkies, police radios, and other two-way radio systems use half duplex communication. “Full duplex” enables both devices to send information simultaneously. For example, telephones and video-conferencing systems use full-duplex communication.

**E**

## Ethernet

A system for connecting a number of computer systems to form a local area network (LAN), with protocols to control the passing of information and to avoid simultaneous transmission by two or more systems. Supports data transfer rates of 10, 100, 1,000, and 10,000 megabits per second (Mb/s). 10, 100, and 1,000Mb/s networks are often referred to as 10BASE-T, 100BASE-T, and 1000BASE-T, respectively. 10,000Mb/s networks are usually referred to as 10Gb Ethernet or 10GbE.

## event log

A record of system events that tracks informational, warning, and error events, such as when significant milestones are reached or when errors occur during activity.

## F

### failover

The capability of a system to switch over automatically to a redundant or standby system upon the failure or abnormal termination of the previously active system. In Nexsan Storage Systems, failover describes one RAID controller taking over the host connections and RAID set control of the other RAID Controller when that controller fails.

### Failover Only

A multipathing load balance policy. It allocates a single active path for each target disk, and any additional paths will be passive. When the active path fails, one of the passive paths will become active, and host I/O will continue down this new path. This is the most basic load balancing mode.

### fault-tolerant

Systems that can continue operating when one or more parts fail are said to be “fault-tolerant”. In Nexsan Storage Systems, the term can be applied to two different areas: the individual RAID sets and the storage system as a whole. RAID sets are said to be fault-tolerant when data is preserved even if one or more disks fail. The storage system as a whole is said to be fault-tolerant if the system mode is set to 2-Port Active-Active (2PAA) mode, 4-Port Active-Active (4PAA) mode, or All Ports All LUNs (APAL) mode and volumes are properly mapped.

### FC port

See “Fibre Channel port”.

### FCC

The Federal Communications Commission; the United States federal agency that regulates electromagnetic emissions.

### Fibre Channel

A gigabit (Gb) speed network technology primarily used for storage networking and the current standard connection type for storage

area networks (SANs). Despite its name, Fibre Channel signaling can run on both twisted-pair copper wire and fibre-optic cables.

### Fibre Channel port

Any entity that actively communicates over a Fibre Channel network. Usually implemented in a device such as disk storage or a Fibre Channel switch. Depending on the system, the Fibre Channel ports on Nexsan Storage Systems can support 2Gb/s, 4Gb/s, or 8Gb/s connections.

### Fibre Channel switch

A network switch compatible with the Fibre Channel protocol. Enables the creation of a Fibre Channel network, which is currently the core component of most storage area networks (SANs).

### Fibre Channel topology

The method of connecting Fibre Channel ports together. A point-to-point (P2P) topology connects two devices directly to each other, with no hub or switch in between. A loop topology usually connects two or more devices in a ring, but can also connect two devices directly to each other just like a P2P topology can. A full-fabric topology connects multiple devices in a network.

### firmware

Small, fixed software applications, stored in read-only memory (ROM) or programmable read-only memory (PROM), that internally control various electronic devices. In Nexsan E-Series, SATABeast/SASBeast, and SATABoy/SASBoy storage systems, each RAID controller is loaded with firmware to control its functionality. Occasionally, this firmware must be updated using the Update Firmware page in the graphical user interface.

### full duplex

See “duplex”.

**G**

## gateway

An internetworking system that joins together the different subnets of a network or two networks that use different base protocols. A network gateway can be implemented completely in software, completely in hardware, or as a combination of both.

## Gb

Gigabit. Approximately one billion (1,000,000,000) bits.

## GB

Gigabyte. Approximately one billion (1,000,000,000) bytes. Used to describe the storage capacity of hard disk drives. A gigabyte is usually computed as  $10^9$  (1,000,000,000) bytes, but can also be computed as  $2^{30}$  (1,073,741,824) bytes (often called a “binary gigabyte” and abbreviated GiB).

## Gb/s

Gigabits (Gb) per second. Used to describe the speed of network data transmission.

## GB/s

Gigabytes (GB) per second. Used to describe the speed of network data transmission. 1 GB/s is eight times faster than 1Gb/s.

## GiB

A binary gigabyte, computed as  $2^{30}$  (1,073,741,824) bytes. See “GB”.

## graphical user interface

A type of user interface that enables users to interact with electronic devices using images rather than text commands. Nexsan Storage Systems use a graphical user interface for system configuration.

## GUI

See “graphical user interface”.

**H**

## host

A computer, server, or other device which accesses the volumes in a Nexsan Storage System. The host can be connected to the Nexsan Storage System with a Fibre Channel connection, an iSCSI connection, or a SAS connection.

**I**

## I/O

Input/Output. The communication between an information processing system (such as a computer or a Nexsan Storage System RAID controller), and the outside world (either an operator or another information processing system). Inputs are the signals or data received by the system, and outputs are the signals or data sent from it.

## IP address

Internet Protocol address. A numerical label assigned to each device (such as a computer, printer, or Nexsan Storage System) on a computer network that uses TCP/IP for communication.

## iSCSI

Internet Small Computer System Interface. A transport protocol that provides for the SCSI protocol to be carried over a TCP/IP network.

**K**

## Kb

Kilobit. Approximately one thousand (1,000) bits.

## KB

Kilobyte. Approximately one thousand (1,000) bytes. Used to describe the storage capacity of hard disk drives and the stripe size in RAIDs. A kilobyte is usually computed as  $10^3$  (1,000) bytes, but can also be computed as  $2^{10}$

(1,024) bytes (often called a “binary kilobyte” and abbreviated KiB).

KiB

A binary kilobyte. Computer as  $2^{10}$  (1,024) bytes. See “KB”.

## L

LAN

See “local area network”.

Least Queue Depth

A multipathing load balance policy. It uses all available paths to target disks, and host I/O is sent to the least busy path based on the number of queued host I/Os. The advantage of this load balancing policy is that host I/O will be evenly spread across the available paths, which can result in better performance (depending on configuration).

load balance policy

In multipathing, a set of instructions for the multipathing software to follow in order to ensure that I/O transfers through host paths are optimally routed and that no one path gets overloaded with data.

local area network

A computer network that links devices within a small geographic area, such as a building or group of adjacent buildings.

logical unit

See “volume”.

LUN

Logical Unit Number. An identification scheme for storage disks that supports a small number of logical units. On Nexsan Storage Systems, LUNs are assigned to volumes and are addressed as LUN 0 through 254.

## M

Mb

Megabit. Approximately one million (1,000,000) bits.

MB

Megabyte. Approximately one million (1,000,000) bytes. Used to describe the storage capacity of hard disk drives. A megabyte is usually computed as  $10^6$  (1,000,000) bytes, but can also be computed as  $2^{20}$  (1,048,576) bytes (often called a “binary megabyte” and abbreviated MiB).

Mb/s

Megabits (Mb) per second. Used to describe the speed of network data transmission.

MB/s

Megabytes (MByte) per second. Used to describe the speed of network data transmission. 1 MB/s is eight times faster than 1Mb/s.

MiB

A binary megabyte. Computed as  $2^{20}$  (1,048,576) bytes. See MByte.

multipathing

A means of presenting volumes to a particular host or hosts via redundant data paths. The intent is to maintain I/O in the event of a path failure. Multipathing may also be used to increase performance. If not configured properly, multipathing may lead to data corruption, as an operating system may not inherently realize that the redundant paths are of the same volume and thus could treat them as different volumes.

Multipathing IO (MPIO)

The Multipathing (MPIO) component installs a Nexsan-specific DSM and Nexsan MultiSan IO tool to manage multipathing.

**MultiSan IO**

One of six Nexsan Storage Tools. MultiSan IO enables Nexsan Storage Systems to load-share over multiple host links (using a load balance policy) and to switch paths in the event of problems on the storage area network.

**N****Nexsan Storage Tools**

A suite of tools that, in addition to the graphical user interface, provide ways to manage Nexsan Storage Systems.

**R****RAID**

Redundant Array of Independent Disks. A system using multiple hard drives organized into a single logical unit for the sharing or replication of data in order to increase data integrity, fault-tolerance, and throughput. Also referred to as a RAID set. RAID levels are organized into RAID levels, which describe their architecture and configuration.

**RAID controller**

A hardware device, software program, or combination of the two which manages the physical disk drives in a RAID and presents them as a single logical unit to attached devices. The RAID Controllers in Nexsan Storage Systems are hardware modules. Nexsan RAID Controllers also provide connections for system administration and configuration.

**RAID level**

A numeric indicator of the architecture used by a RAID. RAID levels can be built using any combination of striping, mirroring, and parity. The levels are numbered from 0 through 6. Some RAID levels can also be combined, and these configurations are usually referred to with a two-digit number. For example, RAID 10 = RAID 1 + RAID 0.

**RAID set**

See "RAID".

**read-only memory (ROM)**

A memory chip that stores values but cannot be changed by normal program instructions. Values in read-only memory are nonvolatile; they are retained even when the unit is powered down.

**reboot**

To restart a computer or computerized electronic device. See also system reboot.

**rolling restart**

On Nexsan Storage Systems, a method of rebooting both RAID controllers in a storage system so that host I/O is not interrupted. Rolling restart is only available if the storage system mode is set to 2-Port Active-Active (2PAA) mode, 4-Port Active-Active (4PAA) mode, or All Ports All LUNs (APAL) mode.

**Round-Robin Load Balancing**

A multipathing load balance policy. It switches all the paths to become active and sends host I/O down all of the available paths, alternating between them in a ring fashion. For example, if a target disk has three paths, the first host I/O will be sent to path 1, the next host I/O to path 2, then path 3, then back to path 1 and so on. Depending on the configuration, improved performance may be achieved due to host I/O being spread across all the available paths.

**Round-Robin Load Balancing with Subset**

A multipathing load balance policy. It operates in the same way as Round-Robin Load Balancing except that it is possible to manually configure individual paths as active or passive using the Multipathing IO (MPIO) management utility. This may be useful if, for example, you know some of the paths are routed through a lower-performance switch, and you therefore wish to make these paths passive.

## S

## SAN

See “storage area network”.

## SAS

Serial Attached SCSI. A serial version of the SCSI interface. A point-to-point architecture that uses a disk controller with four or more channels that operate simultaneously. Each full-duplex channel, known as a SAS port, transfers data at 1.5Gb/s, 3Gb/s, or 6Gb/s in each direction. SAS also supports Serial ATA (SATA) drives, which can be mixed with SAS drives in a variety of configurations.

## SATA

Serial Advanced Technology Attachment. A connection standard for fixed and removable hard disk drives.

## SCSI

Small Computer System Interface. A collection of standards and proposed standards for input/output (I/O) communication, primarily intended for connecting storage subsystems or devices to hosts.

## Single Controller Mode

A system mode for Nexsan Storage System. In this mode, only one RAID controller is active, and failure of this controller makes all arrays and volumes inaccessible.

## storage area network

An architecture that provides for attachment of remote computer storage devices to servers in such a way that the devices appear as locally attached to the operating system.

## Storage Manager

One of the Nexsan Storage Tools. Storage Manager provides a common management point for all Nexsan Storage Systems, either in a standalone window or directly integrated into Windows' Computer Management.

## subnet

A subnetwork, or subnet, is a logically visible subdivision of a TCP/IP network. All computers in a subnet have IP addresses with the same prefix. Addresses in the same subnet are reachable without going through a router, and thus can be reached by broadcast.

## system reboot

On Nexsan Storage Systems, a method of rebooting the unit's RAID controllers. With this method, both controllers are restarted at the same time, which can interrupt host I/O. It is therefore recommended that hosts be disconnected or shut down before performing a system reboot.

## T

## TB

Terabyte. Approximately one trillion (1,000,000,000,000) bytes. Used to describe the storage capacity of hard disk drives. A terabyte is usually computed as  $10^{12}$  (1,000,000,000,000) bytes, but can also be computed as  $2^{40}$  (1,099,511,627,776) bytes (often called a “binary terabyte” and abbreviated TiB).

## TCP/IP

Transmission Control Protocol/Internet Protocol. The set of communications protocols used for the Internet and other similar networks. TCP provides reliable delivery of messages between networked computers. IP uses numeric IP addresses to join network segments.

## TiB

A binary terabyte. Computed as  $2^{40}$  (1,099,511,627,776) bytes. See TB.

## V

## volume

An area of usable storage that is addressed as a single unit as if it were a separate,



physical disk drive. Volumes can exist on a single disk drive or on a RAID that spans multiple disk drives.

## W

### Weighted Paths

A multipathing load balance policy. It enables a numeric value to be assigned to each path, known as the “weight”; host I/O is sent via a single path that has the lowest weight. Therefore, I/O will only be active on a single path regardless of the number of active paths. This enables, in effect, a priority to be set for each path, so that when path failures occur, I/O will be resumed on the best remaining active path.

### World Wide Name

A unique identifier which identifies a particular Fibre Channel or SAS target. Nexsan Storage Systems use two kinds of World Wide Names: World Wide Node Names (WWNN), which identify a device; and World Wide Port Names (WWPN), which identify a specific port on the device.

### WWN

See “World Wide Name”.





**Nexsan** — Sunnyvale, CA, USA

1289 Anvilwood Avenue  
Sunnyvale, CA 94089  
United States of America

**Worldwide Web site**

[www.nexsan.com](http://www.nexsan.com)

E-Series/BEAST support:

[https://helper.nexsansupport.com/esr\\_support](https://helper.nexsansupport.com/esr_support)

Copyright © 2010–2022 Nexsan. All rights reserved.

Nexsan® and the Nexsan logo are trademarks or registered trademarks of Nexsan.

All other trademarks and registered trademarks are the property of their respective owners.

**Nexsan** — European Head Office, UK

Units 33–35 Parker Centre, Mansfield Road  
Derby, DE21 4SZ  
United Kingdom

**Contact**

<https://helper.nexsansupport.com/contact>

This product is protected by one or more of the following patents, and other pending patent applications worldwide: United States patents US8,191,841, US8,120,922; United Kingdom patents GB2466535B, GB2467622B, GB2467404B, GB2296798B, GB2297636B

Part Number: P0450148, Rev. C