



NexentaStor

User Guide 4.0.4

Date: June, 2015

Subject: NexentaStor User Guide

Software: NexentaStor

Software Version: 4.0.4

Part Number: 3000-nxs-4.0.4 000057-B

Copyright © 2015 Nexenta Systems™, ALL RIGHTS RESERVED

Notice: No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose, without the express written permission of Nexenta Systems (hereinafter referred to as “Nexenta”).

Nexenta reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. Nexenta products and services only can be ordered under the terms and conditions of Nexenta Systems’ applicable agreements. All of the features described in this document may not be available currently. Refer to the latest product announcement or contact your local Nexenta Systems sales office for information on feature and product availability. This document includes the latest information available at the time of publication.

Nexenta, NexentaStor, NexentaEdge, and NexentaConnect are registered trademarks of Nexenta Systems in the United States and other countries. All other trademarks, service marks, and company names in this document are properties of their respective owners.

Product Versions Applicable to this Documentation:

Product	Versions supported
NexentaStor™	4.0.4

Contents

Preface	xv
1 Introduction	1
About NexentaStor	1
About NexentaStor Components	2
Using Plugins	2
2 Planning Your NexentaStor Environment	3
About Planning Your NexentaStor Environment	3
Analyze Your Storage Requirements	3
About Redundancy Groups	5
About Device Types	6
3 Using NexentaStor	7
Using the Nexenta Management View (NMV)	7
Starting NMV	7
Logging In to NMV	8
Logging Out of NMV	8
Exiting NMV	9
Refreshing the NMV Display	9
Getting Help	9
Using the Nexenta Management Console (NMC)	9
Logging In to NMC	9
Logging Out of Local or Remote NMC	10
Displaying the NMV URL	10
Using Expert Mode	11
Navigating in NMC	11
Showing Available Commands and Options	11
Selecting Multiple Options	13
Displaying NMC Help	14
Reference List of NMC Help Commands	14

About the Default User Accounts	15
Default Account User Names and Passwords	16
Changing the Default Passwords	16
Scripting	17
Rebooting or Powering Off the Appliance	17
Modifying NexentaStor Configuration Settings	18
Changing Settings Manually	18
Changing Settings Using NMV Wizards	18
Setting Up the Mailer	19
Restarting Standard NexentaStor Services	21
Enabling Kerberos Authentication for SSH	22
Restarting the Nexenta Management Server	22
Configuring NexentaStor as Kerberos Client	23
Configuring NexentaStor as Kerberos Client for KDC in Non-Interactive Mode	23
Configuring NexentaStor as Kerberos Client for NFS with Microsoft Active Directory as KDC Server	25
Validating the Keytab	26
4 Managing Volumes	28
About Volume Management	28
Main Tasks in Using Volumes	29
Creating a Volume	29
Failure Domain	33
Creating a Volume in Manual Mode	34
Viewing the Status of Data Volumes	35
Growing a Volume	36
Attaching a New Disk to a Volume	36
Additional Volume Tasks	37
Displaying Volume Properties	37
Editing Volume Properties	38
Detaching a Disk in a Mirror	42
Removing a Device from a Volume	43
Replacing a Device in a Pool	44
Setting a Device in a Pool to Offline	45

Setting a Device in a Pool to Online	.45
Remounting a Volume	.46
Deleting a Volume	.46
Creating an Auto-Scrub Service	.47
Exporting a Data Volume	.47
Importing a Data Volume	.48
Enabling and Disabling Emergency Reserve	.49
5 Managing Disks and JBODs	.51
About Disk Management	.51
About JBODs	.51
Main Tasks in Managing Disks and JBODS	.52
Viewing the Status of a Disk	.52
Assigning a Device Type	.52
Adding a Spare Device to an Existing Volume	.54
Viewing a JBOD State	.55
Viewing JBOD Disks and Sensor States	.57
Configuring IPMI	.62
Additional Disk and JBOD Tasks	.62
Using Blink	.62
Renaming a JBOD	.63
Using iSCSI Targets as NexentaStor Disks	.64
Using iSCSI Targets without Authentication	.64
Using iSCSI Targets with CHAP Authentication	.65
6 Managing Folders	.67
About Folder Management	.67
Sharing Folders	.68
Controlling Access to Shared Folders	.68
About ACLs	.68
Creating a Folder	.69
Viewing the Status of a Folder	.71
Remounting a folder	.71
Creating a Quick-Backup	.72
Sharing Folders Using NFS	.72

Configuring the NFS Server	73
Sharing a Folder Using NFS	74
Sharing Folders Using CIFS	75
Configuring the CIFS Server	75
Sharing a Folder Using CIFS	76
Anonymous User Access Using CIFS	78
Authenticated User Access Using CIFS in Workgroup Mode	79
Sharing Folders Using FTP	79
Configuring the FTP Server	79
Sharing a Folder Using FTP	80
Sharing Folders Using RSYNC	81
Configuring the RSYNC Server	81
Sharing a Folder Using RSYNC	82
Mounting Shared Folders as Windows Drives	83
Managing Access Control Lists	83
Additional Folder Tasks	85
Unsharing a Folder	85
Editing Folder Properties	86
Deleting a Folder	90
Indexing a Folder	90
Searching a Folder	91
Filtering Folders	92
7 Sharing Folders Using CIFS	93
About CIFS	93
Choosing the Mode	93
Workgroup Mode	93
Domain Mode	94
Workgroup Mode Authentication	94
8 Using Windows Active Directory	97
About Using Windows Active Directory	97
Prerequisites	98
Verifying DNS Settings	98

Adding NexentaStor to Active Directory	99
Viewing Service Logs	100
Setting Permissions	100
Obtaining Administrative Privileges	100
Domain Mode Authentication	101
Joining Active Directory	101
About ID Mapping	104
Mapping Windows Users to NexentaStor Users	104
Default CIFS Accounts	106
Adding Members to Groups	107
Removing Members from Groups	107
Modifying the Access Control List.	108
Troubleshooting	109
Initial Troubleshooting Steps	109
General Troubleshooting	110
Verifying SRV Records of Domain Controller	111
Active Directory Integration	111
Unusual User/Group IDs in ACL	112
9 Managing Users	114
About Users and User Permissions	114
About Default NexentaStor Users	114
NMV User Permissions	115
Creating Users and Setting Permissions	116
Setting the Quota for a User Folder	116
Managing Users	117
About Local Appliance User Groups	117
Creating Local User Groups	118
Setting the Quota for a Group	118
Managing Local User Groups	119
About LDAP Users	119
Supported LDAP Servers	120
Configuring the LDAP Server	120
Example of the Account Entry on the LDAP Server	122

	Uploading an SSL Certificate	122
	Updating the SSL Certificate Information	123
10	Managing the SCSI Target	126
	About the NexentaStor SCSI Target	126
	Creating a Target Portal Group	127
	Creating an iSCSI Target	127
	Creating a Target Group	128
	About ZVOL	128
	Creating a ZVOL	129
	Viewing and Changing ZVOL Properties	130
	Destroying a ZVOL	131
	About LUN Mapping	131
	Creating a LUN Mapping	132
	About Secure Authentication	133
	Establishing Unidirectional CHAP Authentication	133
	Establishing Bidirectional CHAP	133
	Creating a Remote Initiator on a NexentaStor Appliance	134
	Configuring as NexentaStor iSCSI Initiator	134
	Creating an Initiator Group	135
11	Managing Snapshots and Checkpoints	136
	About Snapshots and Checkpoints	136
	Setting up a Periodic Snapshot Service	136
	Using Snapshots	137
	Creating a Snapshot	137
	Viewing Snapshot Summary Data	138
	Cloning a Snapshot	138
	Recovering Files From a Snapshot	139
	Rolling Back to a Snapshot	140
	Deleting a Snapshot	140
	Renaming a Snapshot	141
	Using Checkpoints	141
	Creating a Checkpoint	141

Rolling Back to a Checkpoint	142
Viewing Existing Checkpoints	142
Changing the Number of Checkpoints Kept	143
Permanently Saving a Checkpoint	143
Deleting a Checkpoint	143
Saving Configurations	144
Restoring Configurations	144
12 Managing Auto-Services	147
About Auto-Services	147
Replication Auto-Services	148
Comparing Auto-Tier and Auto-Sync	148
Auto-Tier	148
Auto-Sync	150
General Auto-Services	151
Auto-Snap	151
Auto-Scrub	151
Auto-Smart	152
Creating an Auto-Service	152
Viewing an Auto-Service	152
Editing an Auto-Service	153
Starting and Stopping an Auto-Service	154
Enabling and Disabling an Auto-Service	154
Destroying an Auto-Service	155
13 Monitoring NexentaStor Performance	157
About Performance Statistics	157
Viewing Statistics	157
Available Statistics	159
Using Chart Profiles	163
Removing the Custom Chart	163
Changing the Analytics Settings	163
Viewing I/O Performance	163
Using Performance Benchmarks	164
Running Bonnie Benchmarks	164

	Running IOzone Benchmarks	165
	Running Iperf Benchmarks	166
14	Monitoring SNMP Traps	169
	About NexentaStor SNMP Implementation	169
	Using NexentaStor as SNMP Agent	169
	About SNMP Traps	169
	About NexentaStor MIBs	169
	Using the NexentaStor MIBs	170
	Configuring the SNMP Agent	170
	Using NexentaStor as SNMP Manager	171
	Enabling SNMP Manager	171
	Configuring SNMP Trap Daemon	172
	Creating SNMPv3 Users	173
	Configuring Authentication	174
	Configuring SNMP Trap Handlers	176
	Configuring SNMP Forwarding	177
	Removing SNMPv3 Users	178
	Modifying snmptrapd.conf	178
	Additional SNMP Manager Operations	179
15	Managing Nexenta Management Server	180
	About Nexenta Management Server (NMS)	180
	Multi-NMS	180
	Editing NMS Properties	181
	Editing Web UI Properties	186
16	Managing the Network	188
	About Network Management	188
	About Secure Inter-Appliance Access	188
	Configuring Secure Access Using the IP Table	189
	Configuring SSH-Key Based Authentication	189
	Configuring a Network Interface Card	190
	Unconfiguring a Network Interface Card	191
	Editing a Network Interface	191

Creating a Link Aggregation	192
Creating an IPMP Network Interface	193
Creating an IP Alias	193
Configuring a VLAN	194
Establishing the SSH-Binding	195
NexentaStor TCP Ports	197
17 Using NDMP with NexentaStor	198
About NDMP	198
NDMP Performance Considerations	199
Prerequisites	200
Configuring the NDMP Server	200
Stopping and Restarting the NDMP Server	201
Unconfiguring the NDMP Server	201
Three-Way Copying with NDMPcopy	201
Viewing the NDMP Logs	202
Listing the NDMP Sessions	202
Listing the NDMP Devices	203
Monitoring NDMP Server Performance	203
18 Operations and Fault Management	204
About Runners	204
About Fault Triggers	204
Viewing Fault Triggers	205
Viewing Faults	205
Clearing Faults	206
Checking the Fault Management Facility Status	206
Changing the Content in the Notification Email	206
About NM Watchdog	207
About Collectors	209
Viewing Collectors	210
About Reporters	211
Viewing Reporters	211
About the Indexer	212
Viewing Indexers	213

Viewing Active Runners	213
Enabling and Disabling Runners and Triggers	214
Changing Runner Parameters	215
A NMC Command Line Reference	217
Shell Commands Available in NMC	217
NMC Switch Commands	218
Options Commands	219
Recording for Session	219
NexentaStor Support Request	219
DTrace Framework Commands	219
Query Commands	219
Help Commands	219
Create Commands	220
Destroy Commands	220
Run Commands	220
Share Commands	220
SMTP Commands	221
Setup Commands	221
Show	240
B Fault ID Reference	243
About Fault ID References	243
Displaying the NMS Log file	243
About the NMS fault report	244
Fault Descriptions and Properties	245
CPU Utilization Check	245
Hosts Check	246
Runners Check	247
Service Check	248
Volume Check	249
Memory Check	252
License Check	253
NMS Check	254

NMS Autosmartcheck	254
NMS Rebootcheck	255
SES-Check	255
Quota-Check	255
Net-Check	255
Comstar Trigger	256
NMS ZFS Check	256
NMS FMA Check	256
Disk Check	257
Auto-Sync SVC Trigger	258
Index	260

This page intentionally left blank

Preface

This documentation presents information specific to Nexenta products. The information is for reference purposes and is subject to change.

Intended Audience

This documentation is intended for Network Storage Administrators and assumes that you have experience with data storage concepts, such as NAS, SAN, NFS, and ZFS.

Documentation History

The following table lists the released revisions of this documentation.

Product Versions Applicable to this Documentation:

Revision	Date	Description
3000-nxs-4.0.4 FP1-000057-A	June, 2015	GA

Contacting Support

Methods for contacting support:

- Visit the Nexenta [customer portal](#) or [partner portal](#). Log in and browse a knowledge base.
- Using the NexentaStor user interface, NMV (Nexenta Management View):
 - a. Click **Support**.
 - b. Select an action:
 - **Send by email**
Send the support request to the Nexenta support email.
 - **Save to disk**
Saves the support information to the `/var/tmp` directory on the NexentaStor appliance.
 - c. Complete the request form.
 - d. Click **Make Request**.
- Using the NexentaStor command line, NMC (Nexenta Management Console):
 - a. At the command line, type `support`.
 - b. Complete the support wizard.

Comments

Your comments and suggestions to improve this documentation are greatly appreciated. Send any feedback to doc.comments@nexenta.com and include the documentation title, number, and revision. Refer to specific pages, sections, and paragraphs whenever possible.

Introduction

This section includes the following topics:

- [About NexentaStor](#)
- [About NexentaStor Components](#)
- [Using Plugins](#)

About NexentaStor

The Nexenta Storage Appliance, or *NexentaStor*, is a software-based network-attached storage (NAS) or storage-attached network (SAN) solution. NexentaStor supports file and block storage and a variety of advanced storage features, such as replication between various storage systems and virtually unlimited snapshots and file sizes.

NexentaStor delivers richly-featured software on a software appliance. You can use NexentaStor to manage primary storage in NAS/SAN-based deployments, or to manage second-tier storage for backup.

NexentaStor provides the following features:

- Supports file and block storage
- Allows snapshot mirroring (replication)
- Provides block level replication (remote mirroring)
- Supports NFS v3 and v4, rsync, ssh, and zfs send/receive
- Provides in-kernel CIFS stack
- Supports CIFS and NFS transports for tiering and replication
- Manages large storage pools easily
- Supports direct attachment of SCSI, SAS, and SATA disks
- Supports remote connection of disks using iSCSI, FC, and ATA over Ethernet (AoE) protocols
- Supports 10/100/1GBase-T and many 10G Ethernet solutions, as well as 802.3ad Link Aggregation and multipath I/O
- Integrates with Active Directory and LDAP, including UID mapping, netgroups, and X.509 certificate-based client authentication.

NexentaStor is built on the Illumos kernel. For more information about the Illumos project, see <http://www.illumos.org>.

About NexentaStor Components

NexentaStor comprises the following components:

- **Nexenta Management View (NMV)**

Graphical User Interface that allows you to perform most NexentaStor operations. You can access NMV and manage NexentaStor from Windows, UNIX, or Macintosh operating systems.

- **Nexenta Management Console (NMC)**

A command line interface that enables you to perform advanced NexentaStor operations.

NMC provides more functionality than NMV. The following activities are only available in the NMC:

- System upgrades
- Restarting NexentaStor services
- Certain advanced configurations
- Expert mode operations

- **Nexenta Management Server (NMS)**

The Nexenta Management Server is a service that controls all NexentaStor services and runners. It receives and processes requests from NMC and NMV and returns the output.

- **NexentaStor Plugins**

Plugins are optional modules that provide additional NexentaStor functionality, such as continuous replication or high availability.

Using Plugins

A number of plugins extend the NexentaStor functionality. Several plugins are installed with NexentaStor by default. Other plugins are developed by Nexenta and are available for purchase. Additional plugins have been developed by the community and are free.

See Also:

- <https://community.nexenta.com>: free NexentaStor community project.
- www.nexenta.com: Nexenta's official web-site.

Planning Your NexentaStor Environment

This section includes the following topics:

- [About Planning Your NexentaStor Environment](#)
- [About Redundancy Groups](#)
- [About Device Types](#)

About Planning Your NexentaStor Environment

NexentaStor assists you with creating, provisioning, and managing virtual and physical storage devices. Some methods for using NexentaStor are described below.

Analyze Your Storage Requirements

Before you install NexentaStor, thoroughly analyze your storage requirements.

- Disks:
 - Types of storage required
 - Type of disks required for each type of storage: physical, logical, SSD, or HDD
 - Total storage capacity required for each type of storage
 - Number and sizes of disks required to achieve total storage capacity (measure the usable capacity of disks)
 - Location of physical disks and associated hardware and software
- Storage Format:
 - File storage or block storage
 - All vdevs in a volume must use the same storage type
 - Thick or thin provisioning
 - For file storage: record size, folder structure, Users requiring access, ACLs, etc.
 - For block storage: block size, protocol, etc.
 - You cannot change the block size after creation.

- Redundancy and Backup:
 - Redundancy strategy for each volume, including additional disks is required (all vdevs in a volume should use the same redundancy strategy)
 - Backup and snapshot frequency for each file or target
 - SCSI targets and initiators, mappings, and authentication method
- Architecture and Performance:
 - Number of volumes required
 - Number of vdevs in each volume
 - Number and size of disks in each vdev (all vdevs in a volume should have the same usable capacity)
 - Additional disks required for logs, spares, or L2ARC (cache) in each volume
 - Distribution of volumes and virtual machines in the network
 - Minimum network recommendation is 100Mbps Ethernet

Table 2-1: Summary of Tasks

Task	Instructions
1. On the appropriate servers, install and configure the required physical or logical drives and any associated software.	
2. On the ESXi host, install and configure NexentaStor.	NexentaStor Installation Guide
3. In NexentaStor, create volumes.	Chapter 4, Managing Volumes
4. Create vdevs in each volume. If needed, create an L2ARC vdev in each volume.	Chapter 4, Managing Volumes
5. For any volumes that use file-based storage, create the folders, map shares, Users and ACLs. If you are backing up to tape, configure an NDMP server.	Chapter 6, Managing Folders Chapter 9, Managing Users Chapter 17, Using NDMP with NexentaStor
6. For any volumes that use block storage, set up SCSI targets.	Chapter 10, Managing the SCSI Target
7. Set up the replication autoservices for backups and snapshots of volumes, vdevs, or folders.	Chapter 11, Managing Snapshots and Checkpoints Chapter 12, Managing Auto-Services
8. Monitor the performance of NexentaStor and adjust settings as needed.	Chapter 13, Monitoring NexentaStor Performance Chapter 15, Managing Nexenta Management Server

Table 2-1: Summary of Tasks

Task	Instructions
<p>9. Setup the required network infrastructure. Your network infrastructure may include the following components: multiple network interfaces, HA cluster, linked-aggregated network interfaces, VLANs, IP Aliases, etc.</p>	<p>Chapter 16, Managing the Network</p>
<p>10. If errors occur, note the error code and resolve the error.</p>	<p>Chapter 18, Operations and Fault Management</p>
<p>11. As storage needs increase, add additional storage by adding volumes, adding vdevs to existing volumes, or increasing the size of disks in various vdevs.</p>	<p>Chapter 4, Managing Volumes</p>

About Redundancy Groups

A redundancy group is a group of disks, of the same type, that functions collectively as a unit. This concept is similar to the concept of redundant array of independent disks (RAID) technology. Redundancy groups provide reliable data protection by striping data across multiple disks (RAID-Z1) or by storing identical copies of the same data on several disks (mirror). All of the usable data is protected by check data.

NexentaStor provides a software RAID-Z solution. The main advantage of RAID-Z, compared to a hardware RAID, is the copy-on-write mechanism that eliminates the “write hole” existing in RAID-5.

Write hole is a single point of failure which can occur in case of a power loss and leads to data and parity inconsistency. RAID-Z stores data and checksums in different locations, which eliminates the write hole issue.

Nexenta recommends that you always create a redundancy group for the volume.

NexentaStor offers the following redundancy options:

- **None**
All disks are added to one volume. The size of the volume is equal to the size of all the disks.
- **Mirrored**
Data is replicated identically across two or more disks. All disks in a mirror are the same size, and the size of the mirrored volume is the size of one disk. This configuration can withstand disk failures before data corruption.
- **RAID-Z1 (single parity)**
Data and parity are striped across all disks in a volume. This configuration can tolerate one disk failure before data corruption.
- **RAID-Z2 (double parity)**
Data and parity are striped across all disks in a volume. This configuration can tolerate two disk failures before data corruption.

- **RAID-Z3** (triple parity)

Data and parity are striped across all disks in a volume. This configuration can tolerate three disk failures before data corruption.

About Device Types

NexentaStor is based on ZFS technology, which allows hybrid storage volumes. In ZFS, a storage volume can consist of devices dedicated and optimized for special purposes.

NexentaStor can create the following partitions in a storage pool:

- **Cache** (also called read cache)

Separate cache device used as an Adaptive Read Cache Level 2 (L2ARC) device. Nexenta recommends that you set the cache device to a read-optimized SSD.

- **Log**

Separate log device dedicated to log all the events in the system as transactions. Log devices are usually write-optimized SSDs.

- **Pool (or data)**

Group of storage devices that store all your data.

- **Spare**

Disk that is kept as a reserve. The spare is automatically used as a replacement if any disk in the storage pool fails.

Using NexentaStor

This section includes the following topics:

- [Using the Nexenta Management View \(NMV\)](#)
- [Using the Nexenta Management Console \(NMC\)](#)
- [About the Default User Accounts](#)
- [Scripting](#)
- [Rebooting or Powering Off the Appliance](#)
- [Modifying NexentaStor Configuration Settings](#)
- [Restarting Standard NexentaStor Services](#)
- [Configuring NexentaStor as Kerberos Client](#)

Using the Nexenta Management View (NMV)

Nexenta Management View (NMV) is a graphical user interface that enables you to manage NexentaStor in a user-friendly environment.

Starting NMV

NMV supports the following:

- Internet Explorer v7 or later
- Google Chrome
- Mozilla Firefox v3 or later
- Apple Safari

The URL for NMV is a combination of the NexentaStor IP address and the port number. NMV uses port 8457 as the default port. The NMV URL displays in NMC after installation.

Note:

To ensure that NexentaStor appliance works correctly, verify that all extensions that block pop-up windows or flash are disabled. You may also add a rule to the settings that disables pop-up blocking for the NexentaStor appliance.

See Also:

- [Displaying the NMV URL](#)

- ❖ *To start NMV:*

- ◆ In a supported browser, go to:

```
https://<NexentaStor_IP>:8457
```

Example:

If the NexentaStor IP address is 10.3.60.194, the URL is:

```
https://10.3.60.194:8457
```

Note: You can also use an unprotected connection: http.

Logging In to NMV

You can start NMV and perform many activities without logging in. However, there are some actions that require advanced permissions, such as creating a new volume. When an action requires advanced permissions, you must log in to NMV to complete the action.

NexentaStor includes a default `admin` account that has permission to execute all NMV actions.

- ❖ *To log in to NMV, use one of the following actions:*

- In the upper right corner of the NMV screen, click **Login** and enter your User name and password.
- Wait until a popup login dialog displays stating that advanced permissions are required to complete the action. When the dialog opens, enter the User name and password of a User with the appropriate permissions.

See Also:

- [Default Account User Names and Passwords](#)

Logging Out of NMV

You can log out of NMV without exiting NMV. Logging out of NMV restricts some of the actions you can perform.

If you exit NMV without logging out, you remain logged in. The next time you start NMV, you are still logged in.

- ❖ *To log out of NMV:*

- ◆ In the upper right corner of NMV, click **Logout**.

Exiting NMV

If you are logged in when you exit NMV, your login remains active and you are automatically logged in the next time you start NMV.

- ❖ *To exit NMV:*
 - ◆ Close the tab.

Refreshing the NMV Display

Some NMV screens have a **Refresh** button. Click this button to update the display if you suspect the configuration information is out of date, or if you add hardware or virtual hardware. For example, click **Refresh** if you have just added disks and they do not display in the GUI.

Getting Help

In NMV, you can display the online help system that provides detailed information on how to use NexentaStor.

- ❖ *To display help, using NMV:*
 - ◆ Click **Help**.

A secondary window displays the online help. You can now search for any information.

Using the Nexenta Management Console (NMC)

The Nexenta Management Console (NMC) uses a command line interface instead of a graphical user interface. NMC provides more functionality than NMV.

You can also use NMC to record and replay commands across all deployed NexentaStor instances.

See Also:

- [Scripting](#)

Logging In to NMC

You can log in to NMC using an SSH-client, such as [PuTTY](#).

Logging In to NMC Using a Secure Shell Client

You connect to NMC using a Secure Shell (SSH) client such as PuTTY. You can also use the Console tab in the vSphere client if you are using NexentaStor in a vSphere environment.

NexentaStor provides a default `root` User account for NMC that enables you to perform any operations.

❖ *To log in to NMC using a client:*

1. Open a console connection to the server that contains NexentaStor using one of the following methods:
 - Use an SSH client
 - Use the vSphere client and open the console tab, if using a vSphere environment.
2. In the console, log in at the prompt. The NMC prompt displays after you log in.

Example:

To log in as the root User:

```
myhost console login: root
Password: nexenta
```

System response:

```
nmc@myhost :/$
```

Note:

The NMC command line includes the hostname. This document does not include the hostname in command line examples.

See Also:

- [Default Account User Names and Passwords](#)

Logging Out of Local or Remote NMC

Your NMC session stays open until you log out, even if you exit the console connection.

❖ *To log out of a local or remote NMC session, or Group mode:*

- ◆ Type:

```
nmc:/$ exit
```

Displaying the NMV URL

The URL for NMV is displayed on the first page of the main NMC help screen.

❖ *To display the URL for logging into NMV:*

1. Log in to NMC.
2. Display the main NMC help screen, type:

```
nmc:/$ help
```

3. Locate the line for *Management GUI*. This is the URL for logging into NMV.

Using Expert Mode

NMC provides an expert mode that enables you to execute certain advanced commands in bash shell.

Note: This functionality is only available in NMC.

Caution: Use expert mode with caution. You might accidentally corrupt or delete data, or corrupt the NexentaStor configuration in expert mode.

❖ *To switch to expert mode, using NMC:*

◆ Type:

```
nmc:/$ option expert_mode=1
nmc:/$ !bash
```

You can save the new value for the `expert_mode` in the NMC configuration file by adding the `-s` option to this command. This means that the next time you want to enter bash, you only type `!bash` and you do not need to type: `option expert_mode=1`

❖ *To disable expert mode, using NMC:*

◆ Type:

```
nmc:/$ option expert_mode=0
```

❖ *To exit expert mode, using NMC:*

◆ In the bash shell, type:

```
# exit
```

Navigating in NMC

When you are at the command prompt in NMC, type a command and press **Enter** to execute the command. If the command has available options, the options display. Use the right and left arrows to select the appropriate option and press **Enter** to execute the command.

You can exit NMC dialogs, including help files, and return to the command prompt by typing `q` or `exit` and pressing **Enter**.

Showing Available Commands and Options

This section describes how to display the NMC commands and options that are always available.

Displaying All Commands at the Command Prompt

Pressing **TAB** twice at the command prompt displays a list of all of the commands that you can type there. To get the description of the commands type the command, press spacebar and type --help. You can then type or select the command or option that you want to execute and press **ENTER**.

Table 3-1: List of Some NMC Commands

NMC Command	Description
alias	use alias to view, set and unset aliases at runtime.
bzip2	use bzip2 to compress or decompress from standard input to standard output.
cat	copy standard input to standard output.
create	Create any given object: volume, zvol, folder, snapshot, storage service, and so on.
destroy	Destroy any given object: volume, zvol, folder, snapshot, storage service, and so on.
dtrace	Profiles the appliance for performance and provides detailed views of system internal functions for troubleshooting.
gzip	use gzip to compress or uncompress files
help	Display NexentaStor help and prompts.
query	Advanced query and selection.
record	Start/stop NMC recording session.
run	Execute any given runnable object, including auto-services.
setup	Create or destroy any given object; modify any given setting.
share	Share volume, folder or zvol.
show	Display any given object, setting or status.
support	Contact Nexenta technical support.
unshare	Unshare volume, folder or zvol.

Shortcut for Displaying Commands

You can type the first few letters of a command and then press **TAB** twice to display all commands starting with those letters. If only one command starts with those letters, NMC automatically finishes typing the command.

Example:

If you type the letter `s` at the command prompt and press **TAB** twice, NMC displays all commands starting with `s` and then displays `s` at the command prompt again:

```
nmc:/$ s TAB TAB
scp  setup  share  show  support
```

If you type the letters `se` at the command prompt and press **TAB** twice, NMC finishes typing `setup`, because that is the only command starting with `se`:

```
nmc:/$ se TAB TAB
nmc:/$ setup
```

Shortcut for Displaying Options and Properties

If you type the full name of a command and press **TAB** twice, NMC displays all options for that command and then displays the command again.

Example:

If you type `dtrace` and press **TAB** twice, NMC displays all of the `dtrace` options and then displays `dtrace` at the command prompt:

```
nmc:/$ dtrace TAB TAB
IO      locks      misc      report-all  <Enter>    <?>
cpu     memory     network  comstar     <help|-h
```

When you type a command and one or more letters of an option and press **TAB** twice, NMC lists all matching options. If there is only one matching option, NMC automatically finishes typing the option and displays any additional options.

Example:

If you type `dtrace m` and press **TAB** twice, NMC displays all `dtrace` options that begin with the letter `m`:

```
nmc:/$ dtrace m TAB TAB
memory  misc
```

If only one option begins with the letters you typed, NMC automatically completes the option and displays the additional options or properties associated with that option.

```
nmc:/$ dtrace mi TAB TAB
nmc:/$ dtrace misc
dtruss  dvmstat  execsnoop  opensnoop  statsnoop  <?>
```

Selecting Multiple Options

Some commands allow you to select multiple choices. For example, when you create a volume, NexentaStor displays all available devices and enables you to select one or more devices for the volume.

❖ *To select multiple options in a list, using NMC:*

1. Navigate to the first option using the arrow keys and then press **space bar**.
An asterisk (*) displays next to the option indicating that you selected it.
2. Navigate to another option and press the **space bar** again.
Another asterisk displays by the second option, indicating it is also selected.
3. When all options are selected, press **Enter** to execute the command.

Displaying NMC Help

Type `help` at the command line to display a brief summary of help options and open the main NMC help man page. This file contains information about navigating, using commands, fault management, indexing, and other activities.

Note: The NexentaStor User Guide and main online help reference are only available in NMV.

While you are typing a command or option, there are several ways to display help.

❖ *To display NMC help:*

1. Type the command and `?`.
2. Press `Enter`.

Example:

To display a brief description of the `dtrace` command and list each of its options, type the following:

```
nmc:/$ dtrace ?
```

❖ *To display help for a specific command:*

1. Type the full command with options and `-h`.
 2. Press `Enter`.
- You might need to page down to see the entire help file.
3. Type `q` or `exit` to close the help file and return to the command prompt.

Example:

To display a description of the `show inbox` command, type:

```
nmc:/$ show inbox -h
```

Reference List of NMC Help Commands

The following table provides a reference list of NMC help commands.

Table 3-2: Reference List of NMC Help Command

Type at the System Prompt	System Response
<code>nmc:/\$ TAB TAB</code>	Lists all main NMC commands.
<code>nmc:/\$ <partial string> TAB TAB</code>	Lists all commands that start with the string. If only one command starts with that string, list all subcommands or options for that command. If no response occurs, no commands start with that string.

Table 3-2: Reference List of NMC Help Command

Type at the System Prompt	System Response
<code>nmc:/\$ <command> TAB TAB</code>	Lists all operations for a command. If <code>-h</code> is listed as an option, you can type <code><command> -h</code> for help on the command. If no response occurs, NMC does not recognize the entry as part of a command.
<code>nmc:/\$ help</code>	Displays the main NexentaStor NMC help file.
<code>nmc:/\$ help keyword <string></code> or <code>nmc:/\$ help -k <string></code>	Lists all static commands that include that exact string. Adding the <code>-f</code> option runs a full search that displays all dynamic commands that can include the string.
<code>nmc:/\$ help <command></code> or <code>nmc:/\$ <command> usage</code>	Displays the man page for the specified command (not available for all commands).
<code>nmc:/\$ <command> ?</code>	Lists all available operations with a brief description of each (not available for all commands).
<code>nmc:/\$ <command> help</code>	Displays a detailed help screen for the command (not available for all commands).
<code>nmc:/\$ <command with options> -h</code>	Displays help on the given command (not available for all commands).
<code>nmc:/\$ setup usage</code>	Displays combined man pages for the <code>setup</code> command.
<code>nmc:/\$ show usage</code>	Displays combined man pages for the <code>show</code> command.

About the Default User Accounts

NexentaStor provides several predefined Users: root, admin, and guest.

- **root**
Provides full administrative privileges for NMC. This account cannot be deleted.
- **admin**
Provides full administrative privileges for NMV. This account cannot be deleted.

- **guest**

Provides view-only access to data, status, and reports, in NMV, without requiring logging in. However, when you need to perform an action, such as creating a new user, creating a volume, etc. NexentaStor prompts you to login. This account can be deleted.

Note: You cannot log in to NMC with the guest account.

If you log in to NMC with the `admin` account, a UNIX shell opens. Nexenta does not recommend performing any NexentaStor tasks through the UNIX shell. Using the UNIX shell without authorization from your support provider may not be supported and **may void** your license agreement.

Default Account User Names and Passwords

The following table lists the default NexentaStor User names and passwords.

Table 3-3: NexentaStor Default User Names and Passwords

Default User Account	Login	Password
Admin	admin	nexenta
Root	root	nexenta
Guest	guest	guest

Changing the Default Passwords

Immediately after installing NexentaStor, change the passwords of the `admin` and `root` User accounts.

- ❖ *To change the password of the admin and root account, using NMV:*
 1. Click **Settings > Appliance**.
 2. In the **Administration** panel, click **Password**.
 3. In the **Set Appliance Password** window, type and retype the new password for the appropriate User and click **Save**.
- ❖ *To change the password of the admin User or root account, using NMC:*
 1. Type:

```
nmc:/$ setup appliance password
```
 2. Select either the `root` password or the `admin` password, then follow the prompts to change the password.

Scripting

You can create scripts in NMC. These scripts can run on a schedule or on demand.

- ❖ *For more information about scripting, using NMC:*

- ◆ Type:

```
nmc:/$ create script-runner -h
```

- ❖ *For information about recording, using NMC:*

- ◆ Type:

```
nmc:/$ record -h
```

Rebooting or Powering Off the Appliance

You can power off or reboot (restart) NexentaStor at any time.

- When you power off NexentaStor, you need to manually start it.
- When you reboot NexentaStor, it automatically restarts.

- ❖ *To reboot the appliance, using NMV:*

1. Click **Settings > Appliance**.
2. In the **Maintenance** panel, click **Reboot/Power Off**.
3. In the confirmation dialog, click **OK**.

- ❖ *To reboot the appliance, using NMC:*

- ◆ Type:

```
nmc:/$ reboot
```

- ❖ *To power off the appliance, using NMC:*

- ◆ Type:

```
nmc:/$ poweroff
```

Note:

While NexentaStor is powered off, all shared folders become unavailable and all auto-services stop. The shared folders become available and the auto-services restart when NexentaStor restarts.

Modifying NexentaStor Configuration Settings

After installation and initial configuration, you can change basic NexentaStor settings as needed.

Changing Settings Manually

You can change a number of NexentaStor settings manually.

- ❖ *To change basic NexentaStor settings, using NMV:*
 1. Click **Settings > Appliance**.
 2. In the **Basic Settings** panel, click the name of the appropriate setting.
 3. Modify the setting and click **Save**.
- ❖ *To change basic NexentaStor settings, using NMC:*
 1. Type:

```
nmc:/$ setup appliance
```
 2. Choose the required parameter from the list of available options and follow the onscreen instructions to modify the value.

Changing Settings Using NMV Wizards

NMV contains two wizards (configuration tools) that allow you to configure certain settings or perform common actions after installation.

- Wizard 1 allows you to set or modify the basic configuration, administrative accounts password, and notification system (mailer).
- Wizard 2 allows you to set or modify the network interface and iSCSI initiator; add, export, delete, and import volumes; and create, delete, share, and index folders.

The wizards are not available using NMC.

- ❖ *To run a wizard, using NMV:*
 1. Click **Settings > Appliance**.
 2. In the **Wizards** panel, click **Wizard 1** or **Wizard 2**. The appropriate wizard opens.
 3. In the wizard, make any necessary changes in each screen.
 4. In the last step, click **Save**.

See Also:

- *NexentaStor Installation Guide*

Setting Up the Mailer

Many NexentaStor tasks, such as system failure notification, require that you properly configure the mailer. Normally, you configure the mailer during the installation using Wizard 1. If you did not configure the mailer during the installation, or if you need to change the mailer settings, use the following procedure.

❖ *To set up the mailer, using NMV:*

1. Click **Settings > Appliance**.
2. In the **Administration** panel, click **Mailer**.
3. In the **Set SMTP Mail Server Account** screen, enter the appropriate information for your SMTP server, and the email addresses of Users to receive notifications.
4. If desired, click **Test**.
NexentaStor immediately sends a test email to the address in the **E-Mail Addresses** field.
5. Click **Save**.

❖ *To set up the mailer, using NMC:*

1. Type:

```
nmc:/$ setup appliance mailer show
```
2. Review the value for each property.
3. For each property that needs modification, type the following command and modify the property as needed:

```
nmc:/$ setup appliance mailer <property>
```
4. If desired, test the mailer configuration by typing the following command to immediately send an email to the Users specified in `smtp_addresses`:

```
nmc:/$ setup appliance mailer verify
```

Table 3-4: Options for Modification

Property Name in NMV	Property Name in NMC	Description
SMTP server	<code>smtp_server</code>	SMTP server hostname or IP address.
SMTP User	<code>smtp_user</code>	<ul style="list-style-type: none"> • Name that you use to access your e-mail. • Login name for this SMTP server.
SMTP Password	<code>smtp_password</code>	<ul style="list-style-type: none"> • Password that you use to access your e-mail. • Password for the SMTP server login.
SMTP Send Timeout	<code>smtp_timeout</code>	Length of time that the mailer attempts to send a message before timing out

Table 3-4: Options for Modification

Property Name in NMV	Property Name in NMC	Description
SMTP Authentication	smtp_auth	Method of SMTP authentication that your mail server uses. The options are: <ul style="list-style-type: none"> Plain (default value) SSL TSL
E-mail Addresses	smtp_addresses	One or more existing email addresses to receive all notifications generated by NexentaStor. Separate multiple email addresses with a semicolon (;). If you do not specify an email address, NexentaStor sends all notifications to the email address specified in smtp_to.
E-mail Addresses for faults	smtp_addresses_faults	(optional) One or more existing email addresses to receive all fault notifications generated by NexentaStor. Separate multiple email addresses with a semicolon (;). If you do not specify an email address, NexentaStor sends all notifications to the email address specified in smtp_to.
E-mail Addresses for statistics	smtp_addresses_stats	One or more existing email addresses to receive all statistics and test results generated by NexentaStor. Separate multiple email addresses with a semicolon (;). If you do not specify an email address, NexentaStor sends all notifications to the email address specified in smtp_to.
From E-mail Address	smtp_from	Email address that displays as the sender for all email sent from NexentaStor.
	smtp_to	(optional) Sends message to software support or vendor support. Enter valid TO email address for tech support request. (NMC only)
	smtp_cc	(optional) Sends message to hardware vendor support.
	smtp_keep_time	Represents the number of days to keep the old attachment.
	smtp_max_size	Represents the size in KB.

Restarting Standard NexentaStor Services

NexentaStor contains the services listed in the following table.

Table 3-5: NexentaStor Services

Service Name	Description
cifs-server	Provides shared access to files, printers, serial ports, and miscellaneous communications between nodes on a network.
ftp-server	File transfer protocol. Provides storage and access for sharing large files.
iscsi-initiator	Allows NXS to see the luns from other storage systems or appliances on its local NXS appliance.
iscsi-target	Provides for sharing of NXS appliance volumes through the iSCSI protocol.
ldap-client	Provides method for accessing and maintaining distributed directory information such as a list of Users.
ndmp-server	Provides means to transport data between NAS devices and backup devices.
nfs-client	Allows client to access files over a network through NFS protocol.
nfs-server	File sharing NFS server.
nmdtrace-server	Provides troubleshooting capabilities.
nmv-server	NXS appliance that provides a GUI interface.
ntp-client	Provides method for clock synchronization with NTP server.
rr-daemon	Provides a mechanism through ZFS to convert the file system data to a stream which can be transferred over the network and converted back to a ZFS file system.
rsync-server	Synchronizes files and directories from one location to another.
snmp-agent	Allows management of Simple Network Management Protocol objects defined by different processes through a single master agent.
ssh-server	Uses the secure shell protocol to accept secure connections from remote computers.
syslog-daemon	Separates the software that generates messages from the system that stores them and the software that reports and analyzes them.

Restart these services in case of a system failure.

Note: This functionality is only available using NMC.

❖ *To restart network services, using NMC:*

1. Type:

```
nmc:/$ setup network service
```

2. Choose the required service from the list of available options.
3. Select `restart` from the list of available options.

Enabling Kerberos Authentication for SSH

NexentaStor supports Kerberos authentication for ssh daemon. The default setting for Kerberos is disabled for the security reasons. However, you may want to enable Kerberos authentication according to your environment requirements.

❖ *To enable Kerberos authentication for ssh, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the SSH server pane, click **Configure**.
3. In the **GSSAPIAuthentication** field, select **yes**.
4. In the **GSSAPIKeyExchange** field, select **yes**.
5. Click **Apply**.

❖ *To enable Kerberos authentication for ssh, using NMC:*

1. Type:

```
nmc:/$ setup network service ssh-server configure
```

System response:

```
GSSAPIAuthentication :
yes no
```

2. Select **yes**.

System response:

```
GSSAPIKeyExchange :
yes no
```

3. Select **yes**.
4. Leave the default settings for other parameters.

Restarting the Nexenta Management Server

You can restart the Nexenta Management Server (NMS) in case of a system failure.

Note: This functionality is only available using NMC.

❖ *To restart NMS, using NMC:*

1. Type:

```
nmc:/$ setup appliance nms restart
```

2. When prompted, type “y” to confirm the restart.

3. When prompted, press any key to log in again. You need to start a new SSH session.

Configuring NexentaStor as Kerberos Client

You can configure NexentaStor as a Kerberos client to use it as an authentication application on behalf of users who need access to a resource. To accomplish this, set up a secure session between the client and the service hosting the resource, once the user is recognized as trusted.

To configure NexentaStor as Kerberos client, configure the following components:

- The master KDC.

When you set up Kerberos in a production environment, Nexenta recommends you to have multiple slave KDCs with a master KDC for continued availability of the Kerberized services. Nexenta recommends that you use the master KDC server to run exclusively the KDC service. Do not use this server for any other services.

- DNS server for forward and reverse lookups for master KDC, NexentaStor, and other clients.
- Synchronisation protocol such as NTP, so that the hosts time synchronizes with the chosen KDC time.

Warning:

If the local host time does not synchronize with KDC time, configuration may fail. For Kerberos to function correctly, the time difference between the local host and KDC should not exceed five minutes. If the time does not synchronize, Kerberos clients cannot authenticate to the server. The KDC server might also be vulnerable to security threats.

Configuring NexentaStor as Kerberos Client for KDC in Non-Interactive Mode

You can configure NexentaStor as Kerberos client with root permission either interactively or non-interactively. In the non-interactive mode, you may provide the inputs using the command-line options, or a combination of profile and command-line options. You can use the interactive mode when the command is run without options.

❖ *To configure the NexentaStor NFS server as Kerberos client in the non-interactive mode, using NMC:*

1. Type:

```
nmc:/$ setup appliance kerberos-client configure
```

2. Add an option(s) from the following table.

Table 3-6: Available Options

Option	Description
-n	Configures the machine for Kerberised NFS.
-R <realm>	Configures Kerberos realm.
-k <kdc, kdc_list>	Lists KDC host names. If you do not specify the master KDC, the command uses the first KDC server in the kdc_list.
-a <adminuser>	Defines the Kerberos administrative user.
-c <filepath>	Identifies the path name to the <code>configuration</code> master file on a remote host. If the file does not exist, create the file with appropriate configuration or copy the file from the working Kerberos client.
-d <dnsarg>	Lists the DNS lookup options that should be used in the <code>configuration</code> file. Valid entries include: <code>none</code> , <code>dns_lookup_kdc</code> , <code>dns_lookup_realm</code> , and <code>dns_fallback</code> .
-f <fqdn_list>	Creates a service principal entry associated with each of the listed FQDNs.
-H <logical_host_name>	Identifies that the Kerberos client is a node in a cluster. You must copy the resulting <code>configuration</code> file to the other servers in the cluster.
-m <master_kdc>	Identifies the master KDC that Kerberos client should use.
-p <profile>	You may specify the password only in the profile section. For security reasons, you cannot set the password in the command line. For non-interactive mode: <ol style="list-style-type: none"> 1. Set password in the profile: <code>ADMINPASS <password></code>. 2. Set other required options either in command line or in the profile or in both. 3. Set the <code>-y</code> option only in the command line. Any other <code>PARAM</code> entry is considered invalid and is ignored. The command line options override the <code>PARAM</code> values listed in the profile.
-s <pam_service>	Updates the <code>PAM</code> <code>configuration</code> file to include in the existing authentication stacks.
-T <kdc_vendor>	Configures the Kerberos client to associate with a third party server. List of available KDC server vendors: <code>mit</code> — MIT KDC server. <code>ms_ad</code> — Microsoft Active Directory - You need to know the administrative password to associate with the AD server. <code>heimdal</code> — Heimdal KDC server. <code>shishi</code> — Shishi KDC server.
-N <ad_dns_server>	Identifies the DNS server for Microsoft Active Directory.

Option	Description
-A <ad_domain_name>	Specifies the Microsoft Active Directory domain name.
-y	Skips confirmation dialog by responding Yes.

Note: If you provide the same option in profile and in command line, command line option overrides the profile option.

Configuring NexentaStor as Kerberos Client for NFS with Microsoft Active Directory as KDC Server

You can set up NexentaStor as a Kerberos client to use it as an authentication application. Before you start configuring NexentaStor as Kerberos client with Microsoft Active Directory as KDC server, configure the following components:

- DNS server for forward and reverse lookups for master KDC, NexentaStor, and other clients.
- Synchronisation protocol such as NTP, so that the hosts time synchronizes with the selected KDC time.
- Create user account for NFS client.

❖ *To Configure NexentaStor as Kerberos Client with Microsoft Active Directory as KDC Server, using NMC:*

1. Type:

```
nmc:/$ setup appliance kerberos-client configure
```

System response:

```
Is this a client of non-Solaris KDC? (y/n)
```

2. Type `y`.

3. Select `ms-ad` as KDC server vendor.

```
ms-ad mit heimdal shishi
```

4. Enter the Active Directory domain name.

5. Enter the Active Directory adminuser name.

6. Type the password for the Active Directory user name you entered.

The procedure generates the keytab and the configuration file.

7. Optionally, you may check the created keytab for the principal, using the `klist` and `kinit` utilities.

See Also:

- [Validating the Keytab](#)

Using Ktutil

Use the `ktutil` utility when you need to combine two or more keytabs or to add the generated keys to NexentaStor keytab. For example, when you need to add one keytab for CIFS and another for NFS mapped to different accounts, you may use `ktutil` utility.

For the purpose of an example, the following prerequisites are used:

- Realm name: `EXAMPLE.COM`
- NexentaStor FQDN: `nfsc1.example.com`

Note: To combine two or more keytabs using the utility, switch to bash mode as explained in [Using Expert Mode](#).



1. Log in to bash:

```
nmc:/$ option expert_mode=1
nmc:/$ !bash
```

2. To add the keys, type:

```
# ktutil
```

3. To read the Kerberos V5 keytab file into the current `keylist`, type:

```
ktutil: rkt /root/nfsc1keytab
```

4. To write the current `keylist` into the Kerberos V5 `keytab` file, type:

```
ktutil: wkt /etc/krb5.keytab
ktutil: q
```

Validating the Keytab

After you create a keytab, you can validate the generated keytabs using the `kinit` and `klist` utilities. `utility` lists the Kerberos principal and Kerberos tickets held in a credentials cache, or the keys held in a keytab file. For the purpose of an example, the following prerequisites are used:

- Realm name: `EXAMPLE.COM`
- NexentaStor FQDN: `nfsc1.example.com`

Note: To validate the keytab using the utility, switch to bash mode as explained in [Using Expert Mode](#).

- ❖ *To validate the keytab, using NMC:*

1. Log in to bash:

```
nmc:/$ option expert_mode=1
nmc:/$ !bash
```

2. View the list of keys located in the keytab by typing:

```
# klist -k

KVNO Principal
4 nfs/nfsc1.example.com@EXAMPLE.COM
4 nfs/nfsc1.example.com@EXAMPLE.COM
4 nfs/nfsc1.example.com@EXAMPLE.COM
4 nfs/nfsc1.example.com@EXAMPLE.COM
4 host/nfsc1.example.com@EXAMPLE.COM
4 host/nfsc1.example.com@EXAMPLE.COM
4 host/nfsc1.example.com@EXAMPLE.COM
4 host/nfsc1.example.com@EXAMPLE.COM
4 host/nfsc1.example.com@EXAMPLE.COM
```

3. Validate the keytab file by typing:

```
# kinit -k <name_of_principal>
```

The `kinit` utility loads the principal ticket from Kerberos server to the ticket cache file on the NexentaStor appliance. option gets the key for the Kerberos server from local keytab file without prompting for a password.

Example:

```
# kinit -k nfs/nfsc1.example.com@EXAMPLE.COM
```

4. Verify that the principal ticket exists in the ticket cache by typing:

```
# klist

Ticket cache: FILE:/tmp/krb5cc_0
Keytab name: FILE/etc/krb5/krb5.keytab

Ticket                               Service Principal
Starting 06/26/2014 16:42:13          krbtgt/
Expires 06/27/2014 02:42:13          EXAMPLE.COM@EXAMPLE.COM
Renew until 06/27/2014 16:42:134
```

Managing Volumes

This section includes the following topics:

- [About Volume Management](#)
- [Main Tasks in Using Volumes](#)
- [Additional Volume Tasks](#)

About Volume Management

NexentaStor enables you to aggregate the available disks in the system into logical data volumes, and then to allocate file or block-based storage from the data volume. The data volume provides a storage pooling capability, so that the file systems, or blocks, can expand without being over-provisioned.

Using a data volume enables you to:

- Create various redundant configurations which can include cache and log devices.
- Create a dataset consisting of multiple disk drives which display as a single unit.

A data volume provides redundancy capabilities similar in concept to the RAID features of other storage systems. Redundancy options are:

- none
- mirrored
- RAID-Z1 (single parity)
- RAID-Z2 (double parity)
- RAID-Z3 (triple parity)

Note: Nexenta recommends that each volume has redundancy.

There are some important differences between standard RAID features and the redundancy of a NexentaStor volume.

- When a data volume consists of multiple redundancy groups, NexentaStor dynamically stripes writes across the groups. Unlike RAID-0 stripes, the disks participating in the NexentaStor writes are dynamically determined and there are no fixed length sequences.
- RAID-1 assumes that both sides of the mirror are equally current and correct. NexentaStor always relies on checksums to determine if data is valid, instead of assuming that devices report errors on the read requests. If a conflict occurs, NexentaStor uses the most recent data with a valid checksum.
- RAID-5 uses multiple I/O operations if the data being written is smaller than the stripe width. With NexentaStor RAID-Z1, all writes are full stripe writes. This helps ensure that data is always consistent on disk, even in the event of power failures.

See Also:

- [About Redundancy Groups](#)

Main Tasks in Using Volumes

The following sections describe the most common tasks in managing volumes.

Creating a Volume

A volume is a virtual storage pool that consists of one or more virtual devices, which consist of block devices. You must create at least one storage volume in order to work with NexentaStor.

Creating a volume requires selecting one or more disks to add to the first vdev in the volume. Growing a volume consists of creating additional vdevs in the volume.

Before you can create a volume, NexentaStor must have virtual or physical disks available.

Note: If you create a volume using NMV, you may want to first assign types of devices for disks.

For more information, see [Assigning a Device Type](#).

❖ *To create a volume, using NMV:*

1. Click **Data Management > Data Sets**.
 2. In the **Volumes** panel, click **Create**.
-

Note: Before you create a volume, verify that the disks that you plan to use for do not contain any old volume labels, i.e. they are not included in any other volume configuration.

3. Create a volume configuration by selecting one of the following volume types:

- **Performance**

This type of volume is dedicated for a heavy IOPS environment, such as cloud storage, databases, virtualization, and so on.

- **Balanced**

Select the balanced type of volume if the average number of read and write IOPS is equal.

- **Capacity**

This type of volume is recommended for the environments where the number of write operations is significantly higher than the number of read operations. For example, a backup/archiving server, a disaster recovery server, and so on.

CREATE NEW VOLUME

Name:

Failure domain:

Profile: Performance
 Balanced
 Capacity

Data disks 4/6

Spare

Cache disks

Log disks

Use SSD disks only (0 available)

[Switch to manual mode](#)

Legend:
■ Data 4GB ■ Parity 4GB
■ Spare 0 ■ Log 0
■ Cache 0

Note: If you have disks that are included in an exported volume, the **Use disks with data** checkbox displays. You can use these disks for the new volume. However, data on these disks will be overwritten and the old volume will be destroyed.

4. If you use a JBOD or chassis, select a Failure domain.
See [Failure Domain](#).
5. Optionally, view the graphical representation of disk layout.
The layout shows how disks are spread across JBODs.
6. If you want to create a volume that consist of SSDs only, select the **Use SSD disks only** checkbox.
7. Click **Create**.
If you have less than four disk drives of the same size available in your system, switch to manual mode.

❖ *To create a volume, using NMC:*

1. Type:

```
nmc:/$ setup volume create
```


The prompts guide you through a series of selections.
2. Enter a volume name at the prompt.
Example:

```
Volume: vol1
```

3. Create a group of devices for the volume by selecting one or more options from the list using the arrow keys. Then press **Enter**.

Example:

Group of devices: (Use SPACEBAR for multiple selection)

c1t1d0

4. Select a redundancy type.

The number and type of options listed vary depending upon the selections in [Step 3](#). Options include: pool, mirror, RAID-Z1, RAID-Z2, and RAID-Z3.

Example:

Group redundancy type:

pool

5. Select a compression algorithm setting for this dataset. Press **Enter** to accept the default, on.
6. Confirm the request to create the volume by typing **y**.

Example:

Create volume 'vol1'? y

System response:

```

volume : vol1
  state : ONLINE
  scan  : none requested
config : NAME      STATE      READ WRITE CKSUM
        vol1      ONLINE      0     0     0
        clt2d0   ONLINE      0     0     0
errors  : No known data errors

NAME  SIZE  ALLOC  FREE  CAP  DEDUP  HEALTH  ALTROOT
vol1  7.94G 103K   7.94G  0%   1.00x  ONLINE  -

Volume 'vol1' created successfully. ...

```

7. Confirm the request to create auto-snap services by typing **Y**.

Example:

Create default auto-snap services for 'vol1'? Y

System response:

```

PROPERTY      VALUE
service       : auto-snap
instance      : vol1-000
folder        : vol1
frequency     : daily at 00:00
status        : offline* since 12:49:18
enabled       : true
state         : idle
comment       : default-service
exclude       :

```

```

keep           : 30
last_replic_time : 0
latest-suffix  :
period        : 1
snap-recursive : 1
time_started   : N/A
trace_level    : 1
uniqid        : #####
log           : /var/svc/log/system-filesystem-zfs-auto-snap:vol1-
              000.log

```

The following table describes volume properties that you can set during the volume creation time. You can modify some of these properties after creation. However, you cannot modify properties such as redundancy type, group of devices, and volume name

Table 4-1: Volume Property Definitions

Property (NMV)	Property (NMC)	Definition
Volume Name	volume	Name of the volume defined by User.
Available disks	Group of devices	List of disks available in the system and not used in other pool configurations. You can choose one or multiple volumes according to the desired redundancy configuration.
Redundancy type	Group Redundancy type	Type of the redundancy configuration. As in standard RAID, this enables you to unite multiple hard drives in a group and manage them as a single storage unit. The redundancy configurations provide: <ul style="list-style-type: none"> • Larger storage device • Fault tolerance • Better performance (in some configurations) See Also: About Redundancy Groups
Description	N/A	Description of the volume (optional parameter).
Deduplication	N/A	A technology of eliminating the redundant copies of data and better usage of storage capacity. Use this parameter with caution, because it utilizes RAM resources intensively.
Compression	compression	Controls the compression algorithm for this dataset.
Autoexpand	N/A	Enables automatic LUN expansion. For example, you grow a volume with an additional disk. When you enable <code>autoexpand</code> , you can access the expanded space without exporting and importing the volume or rebooting the system.
Sync	N/A	Controls synchronous requests.
Force creation	N/A	Manually forces volume creation, even when it is in use.

Failure Domain

The purpose behind the concept of Failure Domains is to ensure that no single device failure or maintenance action impacts the integrity of a vDev or Pool. By default, the Auto Profile pool creation mechanism selects the greatest level of redundancy. The user, with knowledge the underlying physical chassis and service processes is allowed to alter the choice of a failure domain.


Failure domain is a level of redundancy in a high capacity JBOD. The options are:

- Chassis
- JBOD
- Caddy

Depending on the equipment you use, you may have a choice of one or more of these three. One chassis may contain multiple JBODs. For example, SuperMicro SuperChassis 847DE26-R2K02JBOD. While not common, a JBOD may contain multiple caddies, where a caddy is a disk carrier which contains multiple disks. In the case where you have a single JBOD that does not contain caddies, only the JBOD failure domain option is available.

The following is an example of a disk layout when a failure domain *caddy* is selected:

Figure 4-1: Example of Disk Layout



In these chassis each caddy contains two disks. When you create a mirrored pool, NexentaStor selects disks from different available caddies in the chassis. Disks from one mirror will never be selected from the same caddy. Therefore, if one disk in the mirror needs to be replaced, the other disk in the mirror remains available.

The following is an example of a disk layout when a failure domain *chassis* is selected:

Figure 4-2: Example of Disk Layout



In this example, when you create a mirrored volume, NexentaStor selects disks from different chassis. Therefore, if one disk in the mirror goes offline, other remains available.

See Also:

- [Assigning a Device Type](#)

Creating a Volume in Manual Mode

If you have less than four disks of the same type and size available in your system, you can use manual mode to create a volume. Use manual mode with caution.

Note: Use manual mode only if you are familiar with the concept of RAID configuration and hybrid pool. Nexenta recommends that you use graphical mode for volume creation. For more information, see [Creating a Volume](#).

- ❖ *To create a volume in manual mode, using NMV:*
 1. In the **Create New Volume** window, select one or more disks from the **Available Disks** list.
 2. Select a **Redundancy Type**.
 3. Click **Add** to move the disks to the **Final Volume Configuration** list.
 4. Enter a **Name** for the volume.
 5. Select additional options for the volume.
 6. Click **Create Volume**.

Viewing the Status of Data Volumes

You can review the status of all volumes in the system.

❖ *To view the status of data volumes, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Volumes** panel, click **Show**.

The **Summary Information: Volumes** window lists the status of each volume in the system, along with additional information about the volumes.

Table 4-2: Volume Parameters Definitions

Volume Parameter	Definition
Volume	Volume name
Configuration	The devices and types of devices in this pool
Size	Total storage pool size
Allocated	Physically allocated space within the pool
Free	Amount of free space on the data volume
Capacity	Percentage of used storage space in the pool
Dedup Ratio	Correlation of actual storage capacity to deduplicated capacity. This value expresses the effectiveness of deduplication usage.
State	Status of the disk
Volume version	
Grow	Adds additional devices to the pool
Export	Unmounts all datasets in the volume
Delete	Destroys the current volume.

❖ *To view the status of data volumes, using NMC:*

◆ **Type:**

```
nmc:/$ show volume <vol_name> status
```

Example:

```
nmc:/$ show volume vol1 status
volume : vol1
state  : ONLINE
scan   : none requested
config: NAME          STATE      READ WRITE CKSUM
       vol1           ONLINE     0    0    0
       c1t1d0         ONLINE     0    0    0
       logs
       c1t2d0         ONLINE     0    0    0
errors : No known data errors
```

Growing a Volume

The Grow function enables you to change the configuration of the data volume and add additional redundancy groups or devices.

Note: NexentaStor must have unconfigured devices available.


- ❖ *To expand (grow) a volume, using NMV:*
 1. Click **Data Management > Data Sets**.
 2. In the **Summary Information: Volume** window, click '+' in the **Grow** column for the appropriate volume.
 3. In the **Grow Volume** window, select one or more disks from the **Available Disks List** and add them as a required device into the **Final Volume Configuration List**.
- ❖ *To expand a volume, using NMC:*
 1. Type:


```
nmc:/$ setup volume <vol_name> grow
```
 2. Respond to the prompts to define the expansion.

Attaching a New Disk to a Volume

NexentaStor enables you to attach a new disk to an existing redundant or non-redundant volume configuration, to add a mirror configuration to the volume. To accomplish this, you must have a spare disk, in the system, which is not used in any other redundant configuration.

Note: The device you add must be at least as large as the existing device or redundancy group.

- ❖ *To attach a new disk to volume, using NMV:*
 1. Click **Data Management > Data Sets**.
 2. In the **Volumes** panel, click **Show**.
 3. Click the appropriate volume to open the **Volume:<volname>** window.
 4. Click **Attach**  for the appropriate disk.
 5. In the **New Disk** dropdown menu, select the new disk.
 6. Click **Attach**.

Note: If replacing an existing disk, the Attach button toggles to Replace.

- ❖ *To attach a new disk to a volume, using NMC:*
 1. Type:


```
nmc:/$ setup volume <vol_name> attach-lun
```
 2. Choose an existing disk with which to form a mirror configuration.

Example:

Existing LUN:
c1t3d0 c1t1d0 c1t2d0

3. Choose a new disk to attach to an existing disk.

Example:

New LUN to attach:
 c1t3d0

4. Type Y to confirm the attachment of a new disk.

Example:

Attach 'c1t3d0' to form a mirror configuration (or be added into existing mirror, if it exists) alongside 'c1t1d0' in the volume 'vol1'? (y/n) y

System response:

The operation will take some time. Run 'show volume vol1 status' to see the in-progress status.

Additional Volume Tasks

The following sections list additional tasks for managing volumes.

Displaying Volume Properties

You can display the property values for a particular volume. NMC displays more properties for a volume than NMV.

- ❖ *To display volume properties, using NMV:*
 1. Click **Data Management > Data Sets**.
 2. In the **Volumes** panel, click **Show**.
 3. In the **Summary** window, click the name of the data volume.

A number of properties for the volume display.

- ❖ *To display volume properties, using NMC*

- ◆ Type:

```
nmc:/$ show volume <vol_name> property
```

A number of properties for the volume are displayed.

Editing Volume Properties

You can edit some volume properties after volume creation. You can edit more properties using NMC than using NMV.

❖ *To edit volume properties, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Volumes** panel, click **Show**.
3. In the **Summary** window, click the name of the data volume.
4. In the **Configurable Properties** panel, change the property as needed.
5. Click **Save**.

❖ *To edit volume properties, using NMC:*

1. Type:

```
nmc:/$ setup volume <vol_name> property
```

A number of volume properties are listed.

2. Navigate to the appropriate property using the arrow keys and press **Enter**.
3. Type the new value and press **Enter**.

The following table describes volume properties.

Table 4-3: Advanced Volume Properties

Name (NMC)	Name (NMV)	Description
<code>aclinherit</code>	ACL Inheritance	<p>Controls the inheritance, of the ACL settings, by new files and subdirectories from the parent directory. You can change these properties at any time.</p> <p>The options are:</p> <ul style="list-style-type: none"> • <code>discard</code> — do not inherit the ACL entries • <code>noallow</code> — inherit ACL entries only with deny access type • <code>restricted</code> — inherit ACL entries, excluding write_owner, write_acl • <code>passthrough</code> — mode of newly created files are defined by inherited ACL entries • <code>passthrough-x</code> — files are created with a permission to execute, if this permission is defined in file creation mode and inherited ACL.

Table 4-3: Advanced Volume Properties

Name (NMC)	Name (NMV)	Description
<code>aclmode</code>	ACL Mode	Determines how to modify the ACL entry, when the system calls <code>chmod</code> . The options are: <ul style="list-style-type: none"> <code>discard</code> — remove all the ACL entries except the POSIX entries <code>owner@</code>, <code>group@</code> and <code>everyone@</code> <code>groupmask</code> — User permissions cannot be greater than group permissions. <code>passthrough</code> — ACL entries do not change when you change dataset permissions, except for the POSIX entries <code>owner@</code>, <code>group@</code>, <code>everyone@</code>
<code>atime</code>	Access Time	Determines whether to update the access time volume parameter every time you address files in this dataset. Setting this property to off can result in significant performance improvement, but can result in internal services confusion.
<code>autoexpand</code>		Controls automatic volume expansion after growing the volume.
<code>autoreplace</code>		Controls automatic device replacement. The default value is on, where NexentaStor automatically finds and replaces a newly installed device in the same physical location as the previous device. When set to off, every time you replace a device in the volume, you must run <code>setup volume <volume_name> replace-lun</code> to locate the new disk and replace it in the volume.
<code>canmount</code>	Can Mount	When set to on, you cannot mount the dataset with the <code>mount</code> command. Similar to setting <code>mountpoint</code> to <code>none</code> .
<code>checksum</code>	Checksum	Determines whether to use the proper checksum algorithm. If set to on, automatically uses the default checksum algorithm <code>fletcher4</code> .
<code>compression</code>	Compression	Enables the compression algorithm for a dataset.
<code>copies</code>	Number of copies	Determines the number of copies of user data for a dataset.
<code>dedup</code>	Deduplication	If set to on, enables the deletion of redundant data copies, using the storage capacity more effectively.
<code>devices</code>	Allow Devices	Controls whether device files in a file system can be opened.
<code>exec</code>	Allow Execution	Controls whether programs in a file system can be executed. Also, <code>mmap(2)</code> calls with <code>PROT_EXEC</code> are disallowed when set to off.
<code>groupquota</code>	N/A	Allocates the amount of space a group can use in this folder.
<code>logbias</code>	Log Bias	Controls how ZFS optimizes synchronous requests for this dataset. If <code>logbias</code> is set to latency, ZFS uses the pool's separate log devices, if any, to handle the requests at low latency. If <code>logbias</code> is set to throughput, ZFS does not use the pool's separate log devices. Instead, ZFS optimizes synchronous operations for global pool throughput and efficient use of resources. The default value is latency.

Table 4-3: Advanced Volume Properties

Name (NMC)	Name (NMV)	Description
<code>mlslabel</code>	N/A	Provides a sensitivity label that determines if a dataset can be mounted in a Trusted Extensions zone. If the labeled dataset matches the labeled zone, the dataset can be mounted and accessed from the labeled zone. The default value is none. This property can only be modified when Trusted Extensions is enabled and only with the appropriate privilege.
<code>nbmand</code>	N/A	Controls whether the file system can be mounted with <code>nbmand</code> (Non-blocking mandatory) locks. This property is for SMB clients only. Changes to this property only take effect when the file system is unmounted and remounted.
<code>nms:auto-syncmark</code>		NMS will not let you delete any datasets on the source and destination that is labelled with <code>nms:auto-syncmark</code> . <code>nms:auto-syncmark</code> gets created while creating an auto-sync service.
<code>nms:description</code>	Description	Descriptive comment (optional).
<code>nms:hold</code>		Number of user holds on a particular snapshot. If a hold exists on a snapshot, attempts to destroy that snapshot by using the “ <code>zfs destroy</code> ” command will return EBUSY.
<code>primarycache</code>	N/A	Controls what is cached in the primary cache (ARC). Possible values: all , none , and metadata . If set to all, it caches both user data and metadata. If set to none, it does not cache user data nor metadata. If set to metadata, it caches only metadata.
<code>quota</code>	Quota	Limits the amount of disk space a dataset and its descendents can consume. This property enforces a hard limit on the amount of disk space used, including all space consumed by descendents, such as file systems and snapshots. Setting a quota on a descendent of a dataset that already has a quota does not override the ancestor's quota, but rather imposes an additional limit. Quotas cannot be set on volumes, as the volume size property acts as an implicit quota.
<code>readonly</code>	Read-Only	Controls whether a dataset can be modified. When set to on, you cannot make modifications. Property abbreviation: <code>rdonly</code> .
<code>recordsize</code>	Record Size	Specifies a suggested block size for files in a file system. Property abbreviation: <code>recsize</code> .
<code>refquota</code>	Referenced Quota	Sets the amount of disk space that a dataset can consume. This property enforces a hard limit on the amount of space used. This hard limit does not include disk space used by descendents, such as snapshots and clones.

Table 4-3: Advanced Volume Properties

Name (NMC)	Name (NMV)	Description
refreservation	Referenced Reservation	<p>Sets the minimum amount of disk space that is guaranteed to a dataset, not including descendents, such as snapshots and clones. When the amount of disk space used is below this value, the dataset is treated as if it were taking up the amount of space specified by refreservation. The refreservation reservation is accounted for in the parent dataset's disk space used, and counts against the parent dataset's quotas and reservations.</p> <p>If refreservation is set, a snapshot is only allowed if enough free pool space is available outside of this reservation to accommodate the current number of referenced bytes in the dataset.</p> <p>Property abbreviation: refreserv.</p>
reservation	Reservation	<p>Sets the minimum amount of disk space guaranteed to a dataset and its descendents. When the amount of disk space used is below this value, the dataset is treated as if it were using the amount of space specified by its reservation. Reservations are accounted for in the parent dataset's disk space used, and count against the parent dataset's quotas and reservations.</p> <p>Property abbreviation: reserv.</p>
secondarycache	N/A	<p>Controls what is cached in the secondary cache (L2ARC). Possible values: all, none, and metadata. If set to all, it caches both user data and metadata. If set to none, it does not cache user data nor metadata. If set to metadata, it caches only metadata.</p>
setuid	Allow Set-UID	<p>Controls whether the setuid bit is honored in a volume.</p>
snapdir	Snapshot Visibility	<p>Controls whether the <code>.zfs</code> directory is hidden or visible in the root of the file system</p>

Table 4-3: Advanced Volume Properties

Name (NMC)	Name (NMV)	Description
<code>sync</code>	Sync	Determines the synchronous behavior of volume transactions. Possible values: <ul style="list-style-type: none"> • Standard: Ensures that synchronous volume transactions, such as <code>fsync</code>, <code>O_DSYNC</code>, <code>O_SYNC</code>, and so on, are written to the intent log. (default value) • Always: Ensures that every volume transaction is written and flushed to stable storage, by a returning system call. This value carries a significant performance penalty. • Disabled: Disables synchronous requests. It only commits file system transactions to stable storage on the next transaction group commit, which might be after many seconds. This value provides the best performance, with no risk of corrupting the pool. However, this value is very dangerous because ZFS is ignoring the synchronous transaction demands of applications, such as databases or NFS operations. Setting this value on the currently active root or <code>/var</code> file system might result in out-of-spec behavior, application data loss, or increased vulnerability to replay attacks. Only use this value if you fully understand the risks.
<code>userquota</code>		Specifies storage space quota for a particular user.
<code>version</code>	N/A	Identifies the current version of the volume.
<code>vscan</code>	N/A	Controls whether regular files should be scanned for viruses when a file is opened and closed. In addition to enabling this property, a virus scanning service must also be enabled for virus scanning to occur if you have third-party virus scanning software. The default value is off.
<code>xattr</code>	Extended Attributes	Indicates whether extended attributes are enabled (on) or disabled (off) for this file system.

See Also:

- [About Redundancy Groups](#)

Detaching a Disk in a Mirror

You can detach a disk from a mirrored configuration, if you need to use it in another pool, or if the disk is corrupted.

❖ *To detach a disk from a mirrored configuration, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Volumes** panel, click **Show**.
3. Click the name of the appropriate volume to open the **Volume:<vol_name>** panel.
4. Click **Detach** for the appropriate mirrored disk.

5. Confirm the operation.

❖ *To detach a disk from a mirrored configuration, using NMC:*

1. Type:

```
nmc:/$ setup volume <vol_name> detach-lun
```

Example:

```
nmc:/$ setup volume vol1 detach-lun
```

2. Choose a LUN to detach. Navigate with arrow keys and press **Enter**.

Example:

```
LUN to detach:
c1t3d0 c1t2d0
```

3. To confirm the detach-lun operation, press y.

Example:

```
Detach 'c1t3d0' from volume 'vol1'? (y/n) y
```

System response:

```
Done. Run 'show volume vol1 status' to display the volume configuration
and status.
```

Removing a Device from a Volume

You can remove cache, hot spare, and log devices from a data volume. For example, you can remove a device from one redundant configuration in order to use it in another redundant configuration.

After removing a disk from a volume, the disk remains online in the system.

Caution: You cannot remove pool device from a volume.

❖ *To remove a device from a volume, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Volumes** panel, click **Show**.
3. Click the name of the appropriate volume to open the **Volume: <vol_name>** panel.
4. Click **Remove**.
5. Confirm the operation.

❖ *To remove a device from a volume, using NMC:*

1. Type:

```
nmc:/$ setup volume <vol_name> remove-lun
```

Example:

```
nmc:/$ setup volume vol1 remove-lun
```

2. Choose LUN from the list. Navigate with arrow keys.

Example:

LUN to remove:

c1t2d0

3. Type `y` to confirm removing the LUN.

Example:

Remove 'c1t2d0' from volume 'vol1'? (y/n) y

System response:

Done. Run 'show volume vol1 status' to display the volume configuration and status.

See Also:

- [Viewing the Status of a Disk](#)

Replacing a Device in a Pool

You can replace a device in a pool.

- ❖ *To replace a device in a volume, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Volumes** panel, click **Show**.
3. Click the name of the appropriate volume to open the **Volume:<volname>** panel.
4. Click **Replace** for the appropriate device.
5. In the dropdown menu, select the old disk.
6. In the dropdown menu, select the new disk.
7. Click **Replace**.

- ❖ *To replace a device from a volume, using NMC:*

1. Type:

```
nmc:/$ setup volume <vol_name> replace-lun
```

Example:

```
nmc:/$ setup volume vol1 replace-lun
```

2. In the list, choose the existing LUN to be replaced, then press **Enter**.

Example:

```
LUN to replace
c1t2d0 c1t4d0
```

3. In the list, choose the new LUN for the volume, then press **Enter**.
4. Type `y` to confirm replacing the LUN.

Replacing the lun involves re-silvering the disk and can take some time.

- ◆ Type the following NMC command to learn if the resilver is complete or in-progress:

```
nmc:/$ show volume <vol_name> status
```

Setting a Device in a Pool to Offline

NexentaStor enables you to set a device in a pool to offline. This prevents writing to or reading from the device.

- ❖ *To set a device in a volume to offline, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Volumes** panel, click **Show**.
3. Click the name of the appropriate volume to open the **Volume:<vol_name>** panel.
4. Click the **Offline** icon for the appropriate device.
5. Confirm the operation.

- ❖ *To set a device in a volume to offline, using NMC:*

- ◆ Type:

```
nmc:/$ setup volume <vol_name> offline-lun
```

Note:

You can only use the `detach-lun` command to remove a LUN from a mirror configuration.

You can use `remove-lun` to remove hot spares, cache and log devices.

Setting a Device in a Pool to Online

NexentaStor enables you to set a device in a pool to online.

- ❖ *To set a device to online, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Volumes** panel, click **Show**.
3. Click the name of the appropriate volume to open the **Volume:<vol_name>** panel.
4. Click **Online** for the appropriate disk.
5. Confirm the operation.

- ❖ *To online a device in a volume, using NMC:*

- ◆ Type:

```
nmc:/$ setup volume <vol_name> online-lun
```

Remounting a Volume

NexentaStor volumes are arranged in a hierarchical tree of objects. However, these volumes can be located on different storage devices. To organize volumes into unified tree of objects, volumes must be mounted. When you create a volume, NexentaStor automatically mounts it. However, in the dataset lifetime, the volume may become unmounted. If it happens, you can manually remount the volume.

Note: When you mount or unmount a volume it may become unavailable for network clients.

This functionality is only available in NMC.

❖ *To remount a volume, using NMC:*

1. Type:

```
nmc:/$ setup volume <volname> remount
```

System response:

```
Are you sure you want to remount <volname>, unmounting and mounting the
folder again may cause a temporary loss of client connections. Proceed
to remount the folder? (y/n)
```

2. Type *y*.

System response:

```
FOLD  PROPERTY  VALUE  SOURCE
<volname> mounted  yes    -
```

Deleting a Volume

You can delete any NexentaStor volume. Destroying the data volume unmounts all of the datasets, exports the data with status Destroyed, and deletes all the associated volume services.

❖ *To destroy a data volume, using NMV:*

1. Click **Data Management > Data Sets**.

2. In the **New Summary Information: Volumes** panel, click .

❖ *To destroy a data volume, using NMC:*

◆ Type:

```
nmc:/$ setup volume <vol_name> destroy
```

Example:

```
nmc:/$ setup volume vol1 destroy
```

System response:

```
Destroy volume 'vol1' and destroy all associated shares (including nfs,
cifs, rsync, and ftp shares (if any)) ? Yes
```

```
volume 'vol1' destroyed
```

Warning:

Use this functionality with caution. Make sure that you delete the right volume. Volume deletion does not immediately destroy the data on the disk. If you deleted the wrong volume, you can import it back and the data remains consistent. However, if you create a new storage volume on the base of this disk, you cannot restore the data from the deleted volume.

Creating an Auto-Scrub Service

NexentaStor enables you to set up an auto-scrubbing service that periodically reads the data blocks and checks for consistency errors. If the pool has a redundant configuration, NexentaStor corrects any errors it finds.

Warning:

Scrubbing is a resource-consuming task. It is preferable to schedule scrubbing during a maintenance window.

❖ *To create an auto-scrub service, using NMV:*

1. Click **Data Management > Auto Services**.
2. In the **Auto-Scrub Services** panel, click **Create**.
3. In the **Create Scrubbing Service** panel, select the volume, setup a scrubbing service, and click **Create Service**.

In the **Show Scrubbing Services** panel, a new auto-scrub service displays.

❖ *To create an auto-scrub service, using NMC:*

◆ **Type:**

```
nmc:/$ setup auto-scrub
```

Example:

```
nmc:/$ setup auto-scrub create
Interval      : weekly
Period       : 1
Day          : Sat
Time        : 3am
```


```
About to schedule 'auto-scrub' service for data volume vol1. Proceed? Yes
```

Exporting a Data Volume

You can export a NexentaStor data volume for future migration, backup, or system upgrade procedures. Generally, the export operation flushes all of the unwritten data from cache to disk and unmounts the volume and all its folders and sub-folders.

Warning:

Nexenta Systems, Inc. strongly recommends that you export a volume before importing it to another NexentaStor appliance. By pulling the active disk from the system without explicitly exporting it, you risk receiving a failed disk on the other system after import.

- ❖ *To export a data volume, using NMV:*
 1. Click **Data Management > Data Sets**.
 2. In the **Summary Information: volumes** panel, click  next to the volume to export. Click **OK**.
- ❖ *To export a data volume, using NMC:*
 - ◆ **Type:**

```
nmc:/$ setup volume <vol_name> export
```

Example:

```
nmc:/$ setup volume vol1 export
```

System response:


```
Export volume 'vol1' and destroy all its rsync, and ftp shares ? Yes
The operation will destroy dependent storage services ('auto-scrub:vol1-000', 'auto-snap:vol1-000', 'auto-snap:vol1-001'). Proceed? Yes
volume 'vol1' exported
```

Importing a Data Volume

You can import a NexentaStor volume from a remote NexentaStor appliance to a current NexentaStor appliance on your local network. To recover properly from any failures or unclean shutdowns, NexentaStor replays any unfinished transactions from the log, whether the import was forced or not.

Warning: If using a separate log device and it becomes unavailable, the import fails. Consider using mirrored log devices to protect your log device against failure.

NexentaStor automatically imports the available volumes with all ZFS settings, folders and sub-folders when the system starts.

- ❖ *To import a volume, using NMV:*
 1. Click **Data Management > Data Sets**.
 2. In the **Summary Information: Volumes** panel, click **Import**.
The list of volumes available for import displays.
 3. In the **Import existing volume** panel, click  opposite the volume to import.
- ❖ *To import a volume, using NMC:*
 - ◆ **Type:**

```
nmc:/$ setup volume import <vol_name>
```

Example:

```
nmc:/$ setup volume import
Volume          : vol1
volume         : vol1
state          : ONLINE
scan           : none requested
```



```

config : NAME          STATE      READ WRITE CKSUM
        vol1           ONLINE     0     0     0
          c1t1d0       ONLINE     0     0     0
            logs
              c1t2d0   ONLINE     0     0     0
errors : No known data errors

```

Enabling and Disabling Emergency Reserve

You can create a reserve on a volume to prepare for potential out-of-memory situations. This is helpful, for example, when a volume freezes and you cannot perform any operations.

The NMS property `autoreserve_space` controls this functionality. This property is enabled by default, but you also must enable the volume-check runner.

Caution: Nexenta Systems, Inc. strongly recommends keeping the `autoreserve_space` property enabled.

NMS creates free-space reserve after the volume is 50% full. In order to alert NMS about the occupied space, enable `volume-check`.

❖ *To enable or disable the emergency reserve for a volume, using NMV:*

1. Click **Settings > Preferences**.
2. In the **Preferences** panel, click **Server**.
3. In the **Customize NMS Properties > Autoreserve_Space** drop-down list, select **Yes** or **No** to create or disable the emergency reserve.

❖ *To enable or disable the emergency reserve for a volume, using NMC:*

◆ **Type:**

```
nmc:/$ setup appliance nms property <autoreserve_space>
```

See Also:

- [Appendix B, Fault ID Reference](#)

This page intentionally left blank

Managing Disks and JBODs

This section includes the following topics:

- [About Disk Management](#)
- [About JBODs](#)
- [Main Tasks in Managing Disks and JBODS](#)
- [Configuring IPMI](#)
- [Additional Disk and JBOD Tasks](#)
- [Using iSCSI Targets as NexentaStor Disks](#)

About Disk Management

NexentaStor provides a graphical view for disk management. You can perform the following actions from within the software:

- Aggregate available disks into data volumes and then allocate file or block-based storage.
- Change the volume configuration.
- Map physical slots to Logical Unit Numbers (LUNs) based on the disk Global Unit Identifier (GUID), and monitor and manage the slot mapping.
- Identify each disk by LED blinking.

About JBODs

Just a Bunch Of Disks, (JBOD), is a non-RAID disk configuration that enables you to expand your storage system with multiple disks as soon as the JBOD connects to the storage. A JBOD does not provide the advantages of fault tolerance and performance improvements of RAID, but it enables you to utilize the entire storage space of the disks.

The NexentaStor JBOD managing interface adds the following features:

- JBOD detection
- Monitoring of all device states of the JBOD
- Physical view of a disk-to-slot mapping
- Possibility to detect and remove the failed device with blinking feature
- E-mail warning notification system

Main Tasks in Managing Disks and JBODS

The following sections list the main tasks for managing disks and JBODS.

Viewing the Status of a Disk

NexentaStor provides live disk statistics that indicate the current resource utilization.

- ❖ To view the available disks, with their statistics, using NMV:

- ◆ Click **Settings > Disks**.

- ❖ To view the available disks, with their statistics, using NMC:

- ◆ Type:

```
nmc:/$ show lun iostat
```

Table 5-1: I/O Statistics

Parameter	Description
Disk (NMV) device (NMC)	Disk or device label
r/s	Reads per second
w/s	Writes per second
kr/s	Kilobytes read per second
kw/s	Kilobytes written per second
wait	Length of the transaction queue
actv	Average number of active transactions
wsvc_t	Average time in waiting queue
asvc_t	Average time of active transaction
%w	Average time where transaction are in queue
%b	Transaction in progress

Assigning a Device Type

Before creating a new volume, you can assign disks to be detected by NexentaStor as data, cache, or log devices. For physical HDDs, you can only assign the *data* type. For physical SSDs, you can assign *log* or *cache* types. If you use virtualized drives, you can assign any type of device.

- ❖ To assign a device type, using NMV:

1. Click **Settings > Disks**.
2. In the **Assignment** column, select the type of device.

3. Click **Change disk assignment**.

Now you can create a volume using the assigned disks.

Adding a Spare Device when Creating a Volume

You can add the hot spare when you create a volume, or any time after that.

- ❖ *To add a spare device when creating a volume, using NMV:*
 1. Select **Data Management > Data Sets**.
 2. In the **Volumes panel**, click **Create**.
 3. In the **Create Volume** window, select the required disk(s) to add to the pool, then click **Add to pool**.
 4. Select one or more disk(s) to add as a spare in the **Available Disks** list and click **Add to spare**.
Selected disk must display in the **Final Volume Configuration** list under spare.
 5. Continue adding the other devices to the volume, if needed.
 6. Enter the volume's name and edit properties, if needed.
 7. Click **Create**.

After successful volume creation, NexentaStor guides you to the **Summary Information: Volumes** web-page where you can review the newly created volume's status information.

- ❖ *To add a spare device, when creating a volume, using NMC:*
 1. Perform the steps 1-3 described in [Creating a Volume](#) section.
 2. Type Y to confirm adding the devices to the volume:
 3. Select one or more options from the list using the arrow keys. Then press **Enter**.

Example:

```
Group of devices: (Use SPACEBAR for multiple selection)
c1t2d0 c1t3d0
```

4. Select **spare** for redundancy type.

Example:

```
Group redundancy type :
pool log cache spare
```

You can continue adding as many devices as needed, and select the required redundancy types for them.

5. Continue the volume creation according to [Creating a Volume](#).

Adding a Hot Spare Device

You can also create a hot spare, after creating a volume.

❖ *To add a spare device when not creating a volume, using NMV:*

1. Select **Data Management > Data Sets**.
2. In the **Volumes** panel, click **Show**.
3. Select a volume
4. In the Hot Spares list, select a disk.
5. Click **Apply**.

See Also:

- [About Redundancy Groups](#)
- [About Device Types](#)
- [Viewing the Status of Data Volumes](#)

Adding a Spare Device to an Existing Volume

❖ *To add a spare device to an existing volume, using NMV:*

1. Select **Data Management > Data Sets**.
2. In the **All Volumes** panel click **+** in the **Grow** column.
3. Select the disk from the **Available Disks** list and click **Add to spare**. Selected disk must display in the **Final Volume Configuration** list.
4. Click **Grow**.

❖ *To add a spare device to an existing volume, using NMC:*

1. Type:

```
nmc:/$ setup volume <volume_name> grow
```

2. Add the additional devices.

Select one or more options from the list using the arrow keys. Then press **Enter**.

```
Group of devices: (Use SPACEBAR for multiple selection)
```

```
c1t1d0
```

3. Select spare redundancy type.

```
Group redundancy type :  
pool log cache spare
```

You can add multiple devices with various redundancy type according to your configuration needs.

4. Type `y` to add additional devices, (optional).
To stop adding the additional devices, type `n`.
5. Select compression algorithm setting for this dataset. Press `Enter` to accept the default, `on`.
6. Confirm the volume grow operation.

Example:

```
Grow volume 'vol1'? (y/n) y
```

System response:

```

volume : vol1
state  : ONLINE
scan   : none requested
config : NAME          STATE      READ WRITE CKSUM
        vol1           ONLINE      0     0     0
        ct1d0          ONLINE      0     0     0
spares
        ct2d0          AVAIL
errors : No known data errors

NAME  SIZE  ALLOC  FREE  CAP  DEDUP  HEALTH  ALTROOT
vol1  7.94G  144K  7.94G  0%  1.00x  ONLINE  -

```

The new spare device is added to `vol1`.

Viewing a JBOD State

You can view the available JBODs after you attach a JBOD to NexentaStor.

- ❖ *To view the JBOD states, using NMV:*
 1. Click **Settings > Disks**.
 2. Click **Disks > JBODs panel > JBODs**.

Note: The JBOD panel only displays in NMV when the JBOD is connected to a NexentaStor appliance.

3. In the **JBOD** panel, review the list of JBODs and the statistics information.

❖ To view the state for one or multiple JBODS in the system, using NMC:

◆ Type:

```
nmc:/$ show jbod
```

System response:

NAME	MODEL	SLOTS	SERIAL	FAILED SENSORS
jbod:1	LSI-DE1600-SAS	5/12	50080e521e49b000	-
jbod:2	LSI-DE1600-SAS	2/12	50080e521e497000	-

Table 5-2: JBOD State Parameters

State	Definition
blink	Enables the LED on the JBOD to physically flash. This helps you to locate the disk in a storage rack.
blinking	Shows status of the blinking task.
busy_slots	Number of busy slots (number of slots that have disks inserted).
failed_sensors	Empty field when in a normal state. In case of any errors on any sensor of a JBOD, this field displays the failed sensor.
model	Model of a JBOD.
name	Name of a JBOD in form jbod:1 given to JBOD automatically by NexentaStor.
picture_types	List of the supported picture formats for this type of JBOD (.txt, .jpg, .svg)
product	Product name. For example: Dell MD300 - is a combination of <vendor>-<product>
rename_model	You can rename a model to define it correctly in the system. For example, if <i>vendor</i> for SuperMicro is incorrectly defined as LSI_CORP, this prop must be 1. After renaming, set it to 0.
serial_number	Serial number of a JBOD.
ses_product	Product name, that a JBOD transfers through SES protocol.
ses_vendor	Vendor's name, that a JBOD transfers using the SES protocol.
slot_ids	Identification number of a slot 1 2 3 4 5 6 7 8 9 10 11 12
slots	Number of slots in a JBOD, in form 'xx/yy', where xx is the number of used slots and yy is the total number of slots in the JBOD.
status	Displays JBOD statuses. The statuses of the JBOD can vary according to JBOD model. Refer to the JBOD documentation for more information.
total_slots	Total number of slots in a JBOD

Table 5-2: JBOD State Parameters

State	Definition
vendor	<p>The name of the vendor after renaming. This value differs from <code>ses_vendor</code>, if you renamed model using the <code>setup jbod <jbod_name> model rename</code> command.</p> <p>Example: you can use this functionality to rename SuperMicro JBODs. These JBODs transfer their vendor and product name as <code>LSI_CORP-SAS2X26</code> instead of real vendor/product names. If you rename it, <code>ses_vendor</code> contains <code>LS_CORP</code> and <code>vendor</code> contains your custom vendor's name, such as <code>SuperMicro</code></p>

Viewing JBOD Disks and Sensor States

JBOD status page displays the list of devices in a JBOD, slotmap image and various statistics information about the devices in a JBOD.

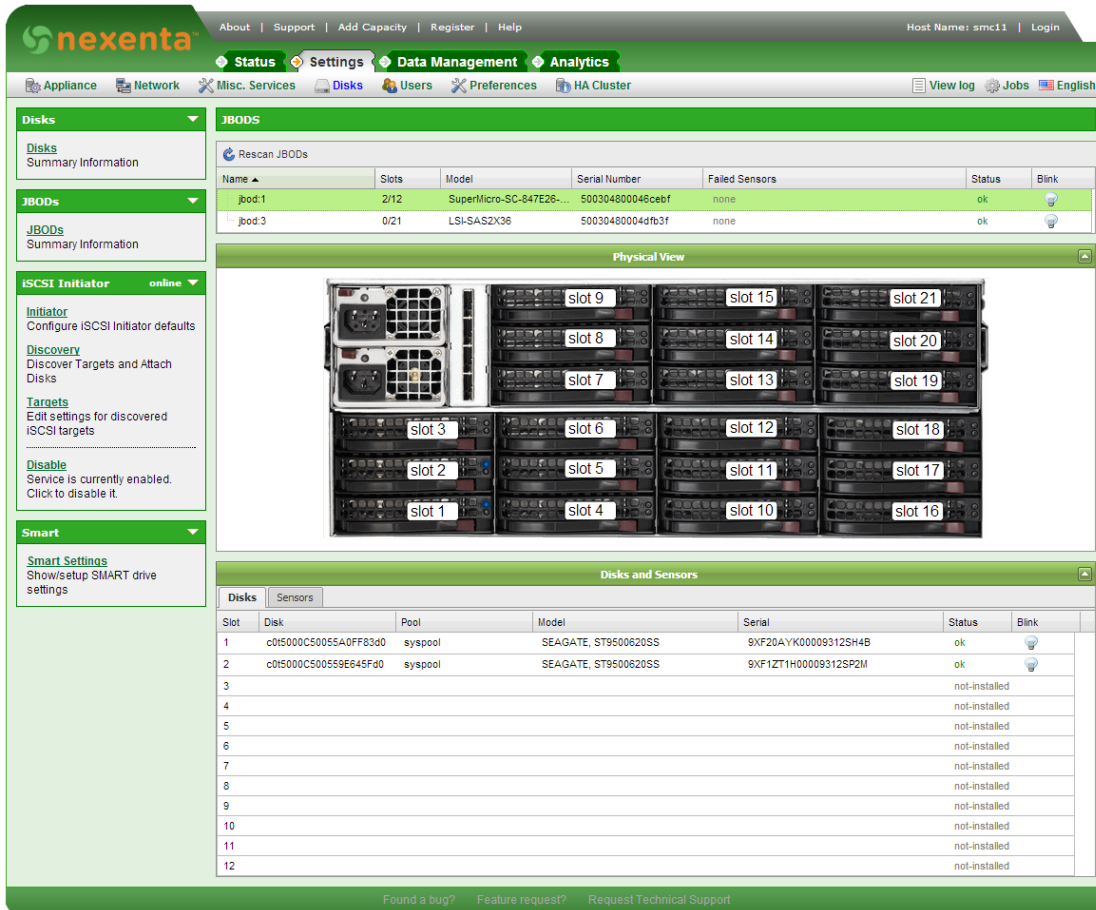
❖ To view JBOD devices, slotmap image, and status, using NMV:

1. Click **Settings > Disks**.
2. Click **Disks > JBODs > JBODs**.
The list of JBODs and the statistics display.
3. Click the JBOD name to see the list of disks and other devices and sensors readings.
4. Expanding the JBOD name also displays the front and the back sides of the corresponding JBOD and the disk-to-slot mapping, as it displays on [JBOD Slot Mapping](#).

Note: NMV displays the front and the back sides of the corresponding JBODS only if it is supported.

5. Click **Disks and Sensors > Disks** or **Sensors** to display the status of disks or sensors.

Figure 5-1: JBOD Slot Mapping



The slot map also displays the state of the disk and the occupied disk slots. The green light indicates that a disk is in a healthy state, the red light indicates that a disk is corrupted, or failed.

Table 5-3: JBOD Disk Panel Parameters

Parameter	Definition
Slot	Slot number according to slot map
Disk	Disk name in NexentaStor. Usually has the following format <i>c0t0d0</i> , where <i>c</i> is the controller number, <i>t</i> is the target number, <i>d</i> is the logical unit number
Pool	Displays volume name to which the device belongs.
Model	Model of a disk
Serial	JBOD serial number
Status	Displays disks statuses. The status of the disk can vary according to JBOD model. Refer to the JBOD documentation for more information.

Table 5-3: JBOD Disk Panel Parameters

Parameter	Definition
Blink	Enables the LED on the disk. It physically flashes so that you can locate the disk in a storage rack

Table 5-4: JBOD Sensor Panel Parameters

Parameter	Definition
Device	Lists the devices in a JBOD, such as: fan, PSU, JBOD.
Sensor	Sensors that helps to monitor JBOD devices state: temperature, voltage, and so on. ¹ If sensor's readings exceed the predefined JBOD thresholds the mailer sends a failure notification to the configured e-mail address. ²
Value	Corresponding value according to sensor type. Example: temperature - C°.
State	Displays device statuses. The statuses of the device can vary according to JBOD model. Refer to the JBOD documentation for more information.

1. For every vendor, each sensor (such as `jbod_temperature1` or `jbod_voltage3`) can indicate a different parameter. Refer to the JBOD technical documentation.
2. You must have the e-mail notification enabled and configured.

❖ *To view JBOD devices and slotmap images, using NMC:*

◆ **Type:**

```
nmc:/$ show jbod <jbod_name> slotmap
```

Example:

```
nmc:/$ show jbod jbod:1 slotmap
```

System response:

```

+-----+
/ /
/ /
+-----+
| [ slot 01 ] [ slot 02 ] [ slot 03 ] [ slot 04 ] | +
| [ slot 05 ] [ slot 06 ] [ slot 07 ] [ slot 08 ] | /
| [ slot 09 ] [ slot 10 ] [ slot 11 ] [ slot 12 ] | /
+-----+
SLOT   LUN                               Pool   Model
1       -                               -       -
2       -                               -       -
3       c4t5000C50025A1050Bd0          -       SEAGATE, ST31000424SS
4       c4t5000C5002109AD93d0          -       SEAGATE, ST3500414SS
5       -                               -       -
6       -                               -       -
7       c4t5000C500212ABCFBd0          -       SEAGATE, ST31000424SS
8       c4t5000C5002109B99Bd0          -       SEAGATE, ST3500414SS
9       -                               -       -
10      -                               -       -
11      -                               -       -
12      c4t5000C5002109B08Bd0          -       SEAGATE, ST3500414SS

```

❖ To view the JBOD disks and sensor parameters, using NMC:

◆ Type:

```
nmc:/$ show jbod <jbod_name> sensors
```

Example:

```
nmc:/$ show jbod jbod:14 sensors
```

System response:

```

ELEMENT  SENSOR      VALUE      STATE
jbod     state           -           ok
slot:1   state           -           ok
slot:2   state           -           ok
slot:3   state           -           ok
slot:4   state           -           ok
slot:5   state           -           ok
slot:6   state           -           ok
slot:7   state           -           ok
slot:8   state           -           ok
slot:9   state           -           ok
slot:10  state           -           ok
slot:11  state           -           ok
slot:12  state           -           ok
slot:13  state           -           ok
slot:14  state           -           ok
slot:15  state           -           ok
slot:16  state           -           ok
slot:17  state           -           ok
slot:18  state           -           ok
slot:19  state           -           ok

```

```

slot:20  state      -          ok
slot:21  state      -          ok
slot:22  state      -          ok
slot:23  state      -          ok
slot:24  state      -          ok
jbod     temperature1  23 C      ok
jbod     temperature2  49 C      ok
jbod     voltage1     5 V       ok
jbod     voltage2     11.99 V    ok
ipmi     System Temp  23 degrees C  ok
ipmi     12VCC        12.26 Volts  ok
ipmi     5VCC         5.03 Volts   ok
ipmi     3.3VCC       3.28 Volts   ok
ipmi     5VSBY        5 Volts      ok
ipmi     3.3VSBY      3.25 Volts   ok
ipmi     1.2VSB       1.27 Volts   ok
ipmi     1.5VSB       1.56 Volts   ok
ipmi     VBAT         3.20 Volts   ok
ipmi     FAN1         7000 RPM     ok
ipmi     FAN2         6700 RPM     ok
ipmi     FAN3         6900 RPM     ok
ipmi     PS1 Status   ok
ipmi     PS2 Status   ok

```

Note: Statistics for sensors available through IPMI is displayed after you configure IPMI.

❖ *To view JBOD failed sensors, using NMC:*

◆ **Type:**

```
nmc:/$ show jbod alerts
```

You can also type `show jbod <jbod_name> alerts` to display the failed sensors for a specific JBOD.

Example:

```
nmc:/$ show jbod :jbod1 alerts
```

System response:

```

jbod:1:
ELEMENT  SENSOR  VALUE  STATE
psu:1    state   -       unknown
psu:2    state   -       unknown

```

Configuring IPMI

NexentaStor provides Intelligent Platform Management Interface (IPMI) protocol support for the JBODs that require IPMI to transfer sensor data. For example, SuperMicro SuperChassis 847DE26-R2K02JBOD and SuperChassis 216BE2C-R920WB.

For the full list of supported JBODs, see *NexentaStor Hardware Supported List*.

NexentaStor collects sensor statistics from the IPMI board on the chassis through the IPMI protocol and displays the statistics on the JBODs' UI page. Sensor statistics are not displayed before you configure IPMI.

Initial setup in NexentaStor must be done for each JBOD.

Before configuring IPMI, obtain the IP address of the chassis power board. See the power board documentation on how to configure the power board IP address in your network.

For Supermicro chassis, see the CSE-PTJBOD-CB3 power board documentation.

❖ *To configure IPMI, using NMV:*

1. Click **Settings > Disks > JBODs**.
The list of JBODs displays.
2. Click **Show IPMI config** in the JBOD panel.
3. In the IPMI column, click the **IPMI icon**.
4. In the **Host name/IP** field, type the IP address of the chassis power board.
5. Type **User name** and **Password**.
6. Click **OK**.

See Also:

- [Viewing JBOD Disks and Sensor States](#)

Additional Disk and JBOD Tasks



The following sections contain additional tasks for managing disks and JBODs.

Using Blink

The Blink feature enables you to make the LED of a JBOD, or of a disk in a JBOD, flash. This functionality can be used to physically locate a failed disk.

❖ *To enable JBOD blinking, using NMV:*

1. Click **Settings > Disks > JBODs**.
The list of **JBODs** and the statistics display.

2. To make an entire JBOD blink:
 - In the JBOD panel, click  next to the required JBOD.
3. To make a disk, or slot, of a JBOD blink:
 1. In the JBOD panel, click on the corresponding JBOD to see the list of disks.
 2. In the list of disks, click on the  next to the required disk.

❖ *To enable JBOD blinking, using NMC:*

◆ Type:

```
nmc:/$ show jbod <jbod:id> blink
```

❖ *To make a disk, or a slot, of a JBOD blink, using NMC:*

◆ Type:

```
nmc:/$ show jbod <jbod:id> lun <disk> blink
```

❖ *To make a fan or psu blink, using NMC:*

◆ Type:

```
nmc:/$ show jbod <jbod:id> <psu/fan> <id> blink
```

After completing these steps, the system administrator can check and perform the required operations with the blinking device, such as removing or replacing the disk. Usually, the LED blinks with orange or blue light.

Note: Blinking does not stop automatically or after a certain period of time. You must click the bulb again once you have finished the operations with a JBOD or disk.

Renaming a JBOD

Sometimes the name of a JBOD is transferred incorrectly to NexentaStor by a SES protocol. If so, rename the model of the JBOD.

❖ *To rename a JBOD, using NMC:*

1. Type:

```
nmc:/$ setup jbod <jbod_name> model rename
```

2. Select the name of a vendor from the list.

If you cannot find the name of the vendor in the list, select **other**, and type the vendor's name.

Note: If you select other for vendor, you do not see any graphic image on NMV or slotmap in NMC.

3. Select the product name from the list.

If you can't find the name of the product in the list, select **other**, and type the product's name.

System response:

```
Renamed successfully.
```

Using iSCSI Targets as NexentaStor Disks

You can provision LUNs from external iSCSI targets as local NexentaStor disks using an iSCSI Initiator. You can attach iSCSI targets without authentication or using CHAP authentication.

Using iSCSI Targets without Authentication

The following describes using two NexentaStor appliances, one configured as an iSCSI target and the second configured as an iSCSI initiator.

- ❖ *Prior to configuring NexentaStor as an Initiator:*
 1. Create an iSCSI target on an external storage or another NexentaStor appliance and share it so that it is available on NexentaStor with an iSCSI Initiator side.
 2. Create a ZVOL on the primary Nexenta appliance, (see [Creating a ZVOL](#)).
 3. Create a target portal group, (see [Creating a Target Portal Group](#)).
 4. Create a target, (see [Creating an iSCSI Target](#)).
 5. Map a previously created ZVOL to the iSCSI target, (see [Creating a LUN Mapping](#)).
- ❖ *To configure the iSCSI Target, using NMV:*
 1. Click **Settings > Disks**.
 2. In the iSCSI Initiator panel, click **Discovery**.
 3. In the iSCSI Discovery Method field, select **Target Address Discovery**.
 4. Type the IP address of the primary iSCSI target.

Note: The iSCSI target can be another Nexenta iSCSI target portal, another NexentaStor appliance, or a third-party storage system.

5. Click **Add Discovery**.
The node IP address displays in the **Configured iSCSI Discovery Methods** field.
6. In the iSCSI Initiator panel, click **Targets**.
The Targets that were just discovered display.
7. In the Disks panel, click **Disks**.
A disk with an extremely long disk name displays.
8. Go to **Data Management > Data Sets** and create a new volume with the iSCSI disks that NexentaStor just discovered.

Now, you can use the target as a regular disk. You can create volumes, folders, share them, etc.

Using iSCSI Targets with CHAP Authentication

When configuring the iSCSI targets to use CHAP authentication, you must first configure the iSCSI target on the primary appliance:

❖ *To configure the iSCSI target:*

1. Click **Data Management > SCSI Target**.
2. In the iSCSI panel, click **Remote Initiators**.
3. Click **Create**.
4. In the Name field, type the Initiator node IQN.
5. In the CHAP User and CHAP secret fields, type appropriate information.
Note that the CHAP secret is 12-16 characters.
6. Click **Create**.
7. In the iSCSI panel, click **Targets**.
8. Click on the target to view its properties.
9. Change the Auth Method to **chap**.

Warning: DO NOT type any data in the CHAP User or CHAP Secret fields.

10. Click **Update**.

Configuring Initiators

After configuring the iSCSI target, you can configure the Initiator.

❖ *To configure the Initiator:*

1. Click **Settings > Disks**.
2. In the iSCSI Initiator panel, click **Initiator**.
3. Change Authentication method to **CHAP**.
4. Enter a CHAP Authentication Name and CHAP Password (12-16 characters).
Use the same name and password that you configured for the target node's Remote Initiator.
5. Click **Save**.
6. In the iSCSI Initiator panel, click **Targets**.
7. In the iSCSI Discovery Method field, select **Target Address Discovery**.
8. Type the IP address of the primary iSCSI target.

Note: The iSCSi target can be another Nexenta iSCSi target portal, another NexentaStor appliance, or a third-party storage system.

9. Click **Add Discovery**.
10. In the Disks panel, click **Disks**.

Check that the new LUN is displays in the list.

Note: You may need to restart (i.e. disable/enable) the iSCSI Initiator service. If something does not work, check the system log on both nodes. Failed login attempts are logged there.

Managing Folders

This section includes the following topics:

- [About Folder Management](#)
- [Sharing Folders](#)
- [Controlling Access to Shared Folders](#)
- [Creating a Folder](#)
- [Viewing the Status of a Folder](#)
- [Remounting a folder](#)
- [Creating a Quick-Backup](#)
- [Sharing Folders Using NFS](#)
- [Sharing Folders Using CIFS](#)
- [Sharing Folders Using FTP](#)
- [Sharing Folders Using RSYNC](#)
- [Mounting Shared Folders as Windows Drives](#)
- [Managing Access Control Lists](#)
- [Additional Folder Tasks](#)

About Folder Management

NexentaStor uses ZFS, which is a standard POSIX compliant dataset. A folder is a manageable storage unit that enables you to organize and share your data over the network. After creating the storage pool, which provides the logical space for creation of the datasets, you can create the hierarchy of multiple folders. The storage pool is always a root directory for ZFS filesystem hierarchy.

NexentaStor unites a number of advanced capabilities to share the folders over the network. You can share NexentaStor folders with various sharing protocols.

NexentaStor folders are managed by multiple properties that enables the User to achieve maximum performance and optimization.

NMV provides a graphical view of the NexentaStor folders, which enables you to:

- See the size, referred and used space of the folder
- Share a folder using various sharing protocols (CIFS, NFS, RSYNC, FTP)

- Manage the Access Control Lists
- Index the folders
- Search for specified folders

See Also:

- [Editing Folder Properties](#)

Sharing Folders

Sharing a folder allows Users to remotely access a folder and its contents.

You can select from several types of sharing protocols. Each protocol handles anonymous and authenticated Users in a different manner.

Note: If you want to mount a NexentaStor folder in Windows, you must share it first in NexentaStor . If you mount it before sharing, you cannot see it in NexentaStor .

See Also:

- [Controlling Access to Shared Folders](#)

Controlling Access to Shared Folders

Most folder sharing protocols allow both anonymous User access and authenticated User access.

- Anonymous User access allows any NexentaStor User to read the contents of a shared folder. Some protocols allow anonymous Users to perform additional actions.
- Authenticated User access requires that a User sign in with a unique User name and password. You can use Access Control Lists (ACLs) to assign access privileges for each User and each folder.

NexentaStor supports local Users and groups along with Users and groups defined in an LDAP-based directory service. NexentaStor also supports native ACLs capable of handling CIFS ACLs, NFSv4 ACLs, and POSIX permissions in the same filesystem. Local Users and groups can override centralized LDAP settings.

About ACLs

NexentaStor supports native extended access control lists (ACLs) that are both CIFS and NFSv4 compliant.

Note: Using the group ACLs is much more efficient than using per-User ACLs. For example, if a new User is added to the Administrators group, he is automatically granted with all the group's permissions.

NexentaStor ACLs are native across ZFS, CIFS, and NFSv4, so they have no conflict in how they are operated on. Generally, you manage ACL through the following tasks:

- Local User or LDAP configuration.
- Definition of per-user or per-group capabilities per volume or folder.
- Overall management of ACLs and ACEs system wide, allowing overriding of end User activity through CIFS/NFS.

NexentaStor contains the following default ACLs that you can use as made or edit as needed:

- Owner
- Group
- Everyone

Note:

NFSv3 relies on POSIX permissions, which are a subset of ZFS extended ACLs. This means that NFSv3 clients only check with the POSIX level permissions.

However, even though POSIX permissions may otherwise grant a permission to a User, it is nullified if the extended ACL on the server is defined and otherwise denies that access.

Creating a Folder

You can create a hierarchy of folders and sub-folders according to your needs. Folders enable you to logically organize the information and share it across the network.

❖ *To create a folder, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Folders** panel, click **Create**.
3. In the **Create New Folder** window, type the name of your new folder in the **Folder name**.
4. Review and edit folder's properties, if needed.
5. Click **Create**.

❖ *To create a folder, using NMC:*

◆ Type one of the following:

```
nmc:/$ setup folder create
```

```
nmc:/$ create folder
```

Example:

```
Folder pathname : vol2/svfol
NAME USED      AVAIL    REFER MOUNTED QUOTA DEDUP COMPRESS
vol2/svfol  31K      7.81G   31K    yes      none  off   on
```

Table 6-1: Folder Term Definitions

Name	Definition
Volume	Folder's root filesystem
Folder name	Custom folder's name
Description	Parameter which contains a description of the folder (optional)
Record Size	Size of the file's block
Log Bias	Manages the synchronous I/Os properly and balances the workload. Logbias properties: Latency — Sends data to a separate log device to proceed with low latency. Throughput — Writes data directly to the pool device at high performance and the write workloads spread.
Deduplication	Eliminates the redundant copies of data and uses storage capacity more efficiently. Use with caution, it utilizes RAM resources intensively
Compression	Enables lzjb compression algorithm
Number of Copies	Number of copies of the dataset to store
Case Sensitivity	Determines if folder can contain directories with unique names, case sensitivity, case insensitivity or mixed perspective. This property is extremely important when you share a folder through both NFS and CIFS protocol. Make sure it is set to "mixed", if you plan to share a folder through both NFS and CIFS. Note that you can specify a case sensitivity only at folder's creation time.
Unicode Only	Allows using only unicode
Sync	Controls synchronous requests

You can set the following parameters only when you create a folder:

Caution:

- Volume
- Folder's name
- Case Sensitivity
- Unicode Only

You cannot change these parameters after creating the folder.

Viewing the Status of a Folder

You can view the status of folders at any time. The folder status displays the available, used and referenced size of all folders.

❖ *To view the status of folders, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Folders** panel, click **Show**.

❖ *To view the status of folders, using NMC:*

◆ Type one of the following:

```
nmc:/$ show folder
```

```
nmc:/$ setup folder show
```

Remounting a folder

NexentaStor folders are arranged in a hierarchical tree of objects. However, these folders can be located on different storage devices. To organize folders into unified tree of objects, folders must be mounted. When you create a folder, NexentaStor automatically mounts it. However, in the dataset lifetime, a folder may become unmounted. If it happens, you can manually remount the folder.

Note: When you mount or unmount a folder it may become unavailable for network clients.

This functionality is only available in NMC.

❖ *To remount a folder, using NMC:*

1. Type:

```
nmc:/$ setup folder <folname> remount
```

System response:

```
Are you sure you want to remount <folname>, unmounting and mounting the
folder again may cause a temporary loss of client connections. Proceed
to remount the folder? (y/n)
```

2. Type `y`.

System response:

```
FOLD      PROPERTY  VALUE    SOURCE
<folname> mounted  yes      -
```

Creating a Quick-Backup

Backing up a folder is easy. If you need to create a backup of a folder, you can run the `quick-backup` command. The command creates an Auto-Sync service with default parameters, executes it once according to schedule, and then changes the status to `maintenance`.

❖ *To create a quick-backup, using NMC:*

1. Type:

```
nmc:/$ setup folder <folname> quick-backup
```

Note: Making quick backup to remote hosts require previously defined ssh-binding. See Setup Network ssh-bind for details.

2. Select a destination host or local host.

3. Select a destination dataset.

System response:

```
Validating parameters... OK
```

```
1 show auto-sync ':test1-archive-001' state and properties
  'show auto-sync :test1-archive-001 -v'
```

```
2 show auto-sync ':test1-archive-001' log
  'show auto-sync :test1-archive-001 log'
```

```
4 show volume 'test1' I/O statistics
  'show volume test1 iostat'
```

Press one of the highlighted keys to make a selection, or press any key to quit

See Also:

- NexentaStor Auto-Sync User Guide
- [Managing Auto-Services](#)

Sharing Folders Using NFS

NFS allows you to share folders on Linux and UNIX operating systems. The folder displays as a local resource.

NexentaStor supports the following versions of NFS:

- NFS v3 (use NFS v3 only for NFSv3 environment, such as MacOSX and VMware ESX Server)
- NFS v4 (recommended)

Sharing the folder using NFS protocol provides the following advantages:

- Shared Storage
- Many clients can connect to NFS share and read/write to this share at the same time.
- Uses filesystem advantages
- Compression, deduplication, snapshots, ACLs, thin provisioning, etc.
- Fast performance
- Uses TCP
- Easy to create and manage. Doesn't require any additional resources.

Configuring the NFS Server

Before sharing a folder using NFS, configure the NFS server according to your IT infrastructure needs.

❖ *To configure NFS server, using NMV:*

1. Click **Data Management > Shares**.
2. In the **NFS server** panel, click **Configure**.
3. In the **Manage NFS server settings** window, select **Service State** to enable the service.
4. Review the other tunable options and change them, if required.
5. Click **Save**.

Table 6-2: NFS Server Tunable Options

NFS Server Option	Description
Service State	Determines the state of the server. Default = on
Server Version	Sets the maximum version for a shared folder by the appliance. Default = 4
Client Version	Sets the maximum version for folders mounted by the appliance. Default = 4
Concurrent NFSD Servers	Controls the number of kernel threads created to handle file service requests. Default = 16
NFSD queue length	Controls the number of queued file service requests sent over a connection-oriented transport (such as TCP/IP). Default = 32
Concurrent LOCKD Servers	Controls how many kernel threads are created to handle file lock requests. Default = 20
LOCKD queue length	Controls the number of queued file lock requests sent over a connection-oriented transport. Default = 32

Sharing a Folder Using NFS

Sharing a folder allows Users to remotely access the folder and its contents. You can share a folder and also edit the folder sharing properties.

❖ *To share a folder using NFS, using NMV:*

1. Click **Data Management > Data Sets**.
2. In **Folders** panel, click **Show**.
3. In **Summary Information: Folders**, select **NFS**.

Selecting a sharing protocol for a parent folder automatically assigns that protocol to all child folders. However, individual child folders can have additional sharing protocols.

Example:

If a parent folder uses the CIFS protocol, both of its child folders must also use CIFS. One of the child folders might also use FTP and NFS, and the other child folder might also use RSYNC.

4. In the confirmation dialog, click **OK**.
5. Click **Edit** for the folder you just shared.
6. Modify folder sharing options as needed, then click **Save**.
7. Mount the folder, as a shared drive, on each Windows machine that accesses the folder.

❖ *To share a folder using NFS, using NMC:*

1. Type one of the following:

```
nmc:/$ setup folder <vol_name/fol_name> share
```

```
nmc:/$ share folder <vol_name/fol_name>
```

Example:

```
nmc:/$ setup folder datavol119/folder1 share
```

2. Select a sharing protocol.
3. Select the appropriate value for each option.
4. Mount the folder as a shared drive on each Windows machine that accesses the folder.

Table 6-3: Sharing Options

NFS Folder Sharing Option	Definition
Auth Type	<p>Determines the authentication type.</p> <ul style="list-style-type: none"> • Auth_sys — Insecure authentication, User name and password transferred transparently. • Auth_none — Null authentication. NFS clients mapped by NFS servers as User nobody. • Auth_des — Diffie Helman public key exchange. • Kerberos 5 — Uses protocol Kerberos 5 before granting access. • Kerberos 5i — Uses Kerberos 5 authentication with checksum verifying. • Kerberos 5p — Uses Kerberos 5 authentication with integrity checksum and shared folder encryption. All traffic is encrypted. This is the most secure authentication type.
Anonymous	Grants access to anonymous Users. If you are using an authentication method, leave this option unchecked.
Anonymous Read-Write	Grants read/write access to anonymous Users.
Read-Write	Lists the Users with Read-write access. Overrides the 'Read-Only' field for the clients specified.
Read-Only	Lists the Users with Read-only access.
Root	List of root Users from the other computer hosts that have root access to this share.
Extra Options	Comma-delimited list of options to control access to the shared resource.
Recursive	Determines whether to share all the child folders.

Sharing Folders Using CIFS

NexentaStor provides kernel and ZFS-integrated CIFS stacks, with native support for Windows Access Control Lists (ACL).

Configuring the CIFS Server

- ❖ *To configure the CIFS server, using NMV:*
 1. Click **Data Management > Shares**.
 2. In the **CIFS Server** panel, click **Configure**.
 3. In the **Manage CIFS Server Settings** window, change any options as needed. Make sure the **Service State** is on (selected).

4. Click **Save**.
 5. In the **CIFS Server** panel, click **Join Workgroup**.
 6. In the **Manage CIFS Server Settings** window, change the workgroup name, if needed.
 7. Click **Save**.
- ❖ *To configure the CIFS server, using NMC:*
1. Type:


```
nmc:/$ setup network service cifs-server configure
```
 2. Follow the prompts to change CIFS server options as needed.
- ❖ *To change the workgroup name:*
1. Type:


```
nmc:/$ setup network service cifs-server join_workgroup
```
 2. Restart the CIFS server by typing:


```
nmc:/$ setup network service cifs-server restart
```

The following table lists the CIFS server options.

Table 6-4: CIFS Server Options

CIFS Server Option	Description
Service State	Turns the CIFS server on and off.
Server String	A descriptive name for the CIFS server that can display in various places.
Anonymous Password	Changes the password for the anonymous User <code>smb</code> . You do not need to know the current password in order to set a new one.
Enable Dynamic DNS Update	Enables a DNS client to register and dynamically update their resource records, with a DNS server, whenever changes occur.
LM Authentication Level	For NexentaStor 3.0.x and later, set this to 4.

Sharing a Folder Using CIFS

When you share a folder the NexentaStor users can remotely access the folder and its contents. You can share a folder and also edit the folder sharing properties.

Note: NexentaStor replaces any slashes (/) in the folder name with underscores (_). Slashes are not allowed in the names of shares.

- ❖ *To share a folder using CIFS, using NMV:*
1. Click **Data Management > Data Sets**.
 2. In the **Folders** panel, click **Show**.
 3. In **Summary Information: Folders**, select CIFS.

Selecting a sharing protocol for a parent folder automatically assigns that protocol to all child folders. However, individual child folders can be shared using additional sharing protocols.

Example:

If a parent folder uses the CIFS protocol, both of its child folders must also use CIFS. One of the child folders might also use FTP and NFS, and the other child folder might also use RSYNC.

4. In the confirmation dialog, click **OK**.
 5. Click the **Edit** link for the folder you just shared.
 6. Modify folder sharing options as needed, including the name, then click **Save**.
 7. Mount the folder as a shared drive on each Windows machine that will access the folder.
- ❖ *To share a folder using CIFS, using NMC:*

1. Type one of the following:

```
nmc:/$ setup folder <vol_name/fol_name> share
```

```
nmc:/$ share folder <vol_name/fol_name>
```

Example:

```
nmc:/$ setup folder datavol119/Folder1 share
```

2. Select a sharing protocol.
3. Select the appropriate value for each option.
4. Mount the folder as a shared drive on each Windows machine that will access the folder.

Table 6-5: CIFS Options

CIFS Folder Sharing Option	Description
Share name	By convention, a folder named vol_name/fol_name becomes a CIFS share named vol_name_fol_name.
Anonymous Read-Write	To allow any NexentaStor User to access the folder, select this option or set the property to true. If you are using any form of authentication for this folder, deselect the option in NMV or set the property to false in NMC.
Restrict Client Host Access	Restrict access for the client.
Guest Access	Minimal rights.
Recursive	Shares the folder and its sub-folders.

Anonymous User Access Using CIFS

Anonymous User access provides any person on the network with full access to a shared folder and all child folders, including the ability to read, write, copy, delete, and execute files. Anonymous Users can also mount a shared folder that has anonymous User access.

Caution: An authenticated User with no access or read-only access to a shared folder can use anonymous access to log on to the folder with full privileges.

When you grant anonymous access to a parent folder, you also grant anonymous access to all existing child folders. However, any new child folders you create do not have anonymous access unless you explicitly turn it on for the new child folders.

You can turn off anonymous access for any folder. If you turn off anonymous access for a parent folder, its child folders maintain the anonymous access unless you turn it off individually for each child folder.

Note: To map a CIFS folder on Windows XP, you must enable guest access for the CIFS folder. All Windows XP users are mapped as guest users.
You can enable guest access in the CIFS folder properties using NMV or NMC.
For more information, see [Sharing a Folder Using CIFS](#).

❖ *To connect to a shared folder with anonymous User access, using Windows:*

1. In Windows, select **Start > Computer**.
2. Click **Map network drive**.
3. In the **Folder** field, type the path to the shared folder.

Example:

```
\\192.168.1.10\data_public
```

Note: You must change all forwardslashes to backslashes. Forward slashes between volume and folder, as well as between folder and subfolders must be changed to underscores.

4. Select **Connect using different credentials**.
5. Click **Finish**.
6. Log on with the following credentials:
 - User name: smb
 - Password (default): nexenta

When you change the anonymous password, all anonymous Users must use the new password.

❖ *To change the anonymous User password, using NMV:*

1. Click **Data Management > Shares**.
2. In the **CIFS Server** panel, click **Configure**.
3. In **Manage CIFS Server Settings**, type the new password.

You do not need to know the old password.

4. Depending on your Windows version, modify the ACL of these directories using the Windows ACL editor.
 1. Right click **Properties**.
 2. Select **Security tab**.

Authenticated User Access Using CIFS in Workgroup Mode

Workgroup mode is the default CIFS authenticated User access mode in NexentaStor. The workgroup mode allows you to grant customized access to specific shared CIFS folders on a per-User basis. Users log in to the shared folders using their NexentaStor User name and password.

Note:	Do not use name-based mapping in workgroup mode. If you do, the mapping daemon (called <code>idmap</code>) tries to search Active Directory to resolve the names, and will most probably fail.
--------------	---

To use workgroup mode authentication:

- Make sure the CIFS Server is configured and online. See [Configuring the CIFS Server](#).
- Create all required NexentaStor Users. See [Creating Users and Setting Permissions](#).
- Share the appropriate folders as described in [Sharing Folders Using CIFS](#), but make sure Anonymous Read-Write is off for each folder (uncheck the box in NMV or set the property to false in NMC).
- Map the appropriate folders as network drives on each User's Windows computer.

See Also:

- [Adding NexentaStor to Active Directory](#)
- [About CIFS](#)

Sharing Folders Using FTP

File Transfer Protocol (FTP) is a standard network protocol used to transfer files from one computer host to another using a TCP-based network. The default FTP port is 21.

Configuring the FTP Server

You can change the server settings to optimize your configuration.

- ❖ *To configure the FTP server, using NMV:*
 1. Click **Data Management > Shares**.
 2. In the FTP Server panel, click **Configure**.

3. Perform the required changes and click **Save**.

Table 6-6: FTP Server: Tunable Options

FTP Server Option	Description
Service State	Enable or disable the service.
Anonymous Access	Enables anonymous access to server.
Anonymous chroot	Specifies the <code>chroot()</code> path for anonymous users.
Deny access to host(s)	Denies access to host(s) that match template.
Banner	Welcome message to Users.
Max Idle Timeout	Maximum idle time in seconds.
Keep alive	Sets the <code>TCP_SO_KEEPALIVE</code> option for control and data sockets. If no or empty, use system default (usually off).
Enable log	Enables logging of file transfers made by FTP users.

Sharing a Folder Using FTP

Sharing a folder allows Users to remotely access the folder and its contents. You can share a folder and edit the folder sharing properties using the following procedure.

- ❖ *To share a folder using FTP, using NMV:*

1. Click **Data Management > Data Sets**.
2. In **Folders** panel, click **Show**.
3. In **Summary Information: Folders**, select FTP.

Selecting a sharing protocol for a parent folder automatically assigns that protocol to all child folders. However, individual child folders can have additional sharing protocols.

Example:

If a parent folder uses the CIFS protocol, both of its child folders must also use CIFS. One of the child folders might also use FTP and NFS, and the other child folder might also use RSYNC.

4. In the confirmation dialog, click **OK**.
5. Click **Edit** for the folder you just shared.
6. Modify folder sharing options as needed, then click **Save**.

- ❖ *To share a folder using FTP, using NMC:*

1. Type one of the following:

```
nmc:/$ setup folder <vol_name/fol_name> share
```

```
nmc:/$ share folder <vol_name/fol_name>
```

Example:

```
nmc:/$ setup folder datavol119/folder1 share
```


2. Select a sharing protocol.
3. Select the appropriate value for each option.

Sharing Folders Using RSYNC

RSYNC is a network protocol for Unix-based and Windows systems which synchronizes files and directories, from one location to another, while minimizing data transfer by sending only the delta instead of the entire file. NexentaStor enables you to share any folder using RSYNC. The default RSYNC port is 873. Depending on your OS, use the corresponding RSYNC client to access the share.

Configuring the RSYNC Server

You can change the server settings to optimize your configuration.

- ❖ *To configure the RSYNC server, using NMV:*
 1. Click **Data Management > Shares**.
 2. In the RSYNC Server panel, click **Configure**.
 3. Perform the required changes and click **Save**.

Table 6-7: RSYNC Tunable Options

RSYNC Server Option	Description
Service State	Enable or disable the service.
Listening Port	You can override the default port the daemon listens to by specifying this value. Default port is 873.
User	User name or ID that files are sent to and from. In combination with the "gid," this option determines what file permissions are available. Default is "nobody".
Group	Group name or ID that files are sent to and from. In combination with the "uid," this option determines what file permissions are available. Default is "nobody".
Max Connections	Allows you to specify the maximum number of simultaneous connections (0 - means no limit). Default is 0.
Read-Only	Determines whether clients can upload files or not.
Apply to all	Resets option to read only for all shares of service to the global value.
Welcome Message	Type in a message.

Sharing a Folder Using RSYNC

Sharing a folder allows Users to remotely access the folder and its contents. You can share a folder and also edit the folder sharing properties using the following procedure.

❖ *To share a folder using RSYNC, using NMV:*

1. Click **Data Management > Data Sets**.
2. In **Folders** panel, click **Show**.
3. In **Summary Information: Folders**, select Rsync.

Selecting a sharing protocol for a parent folder automatically assigns that protocol to all child folders. However, individual child folders can have additional sharing protocols.

Example:

If a parent folder uses the CIFS protocol, both of its child folders must also use CIFS. One of the child folders might also use FTP and NFS, and the other child folder might also use RSYNC.

4. In the confirmation dialog, click **OK**.
5. Click the **Edit** link for the folder you just shared.
6. Modify folder sharing options as needed, then click **Save**.

❖ *To share a folder using RSYNC, using MMC:*

1. Type one of the following:

```
nmc:/$ setup folder <vol_name/fol_name> share
nmc:/$ share folder <vol_name/fol_name>
```

Example:

```
nmc:/$ setup folder datavol119/folder1 share
```

2. Select a sharing protocol.
3. Select the appropriate value for each option.

Table 6-8: RSYNC Configurable Options

RSYNC Folder Sharing Option	Description
Share Name	Name of folder
Share Comment	Descriptive comment.
Use Chroot	True or False. Advantage — Extra protection against possible implementation of security holes. Disadvantages — Requires super-user privileges. Cannot follow symbolic links that are either absolute or outside of the new root path. Complicates the preservation of users and groups by name. Default is false .

Table 6-8: RSYNC Configurable Options

RSYNC Folder Sharing Option	Description
Read-Only	Use Global true or false. Determines whether clients can upload files or not. If true, then attempted uploads fail. If false, then uploads possible if file permissions on the daemon side allow them. Default is <code>use_global</code> .
Write-Only	Determines whether clients can download files or not. If true, then attempted uploads fail. If false, then uploads possible if file permissions on the daemon side allow them. Default is false.
Allowed hosts	List of hosts allowed access.
Denied hosts	List of hosts denied access.

Mounting Shared Folders as Windows Drives

Some folder sharing protocols require you to map shared folders as a shared drive on each Windows machine that will access the folder.

When you map a shared parent folder, all child folders are available on the shared drive.

❖ *To map a shared CIFS folder and its subfolders, using Windows:*

1. On the Windows desktop, right-click **My Computer** and select **Map Network Drive**.
2. In the **Folder** field, enter the NexentaStor hostname or IP address and the full filepath of the shared folder.
3. Click **Finish**.
4. For authenticated access, enter the Windows User's NexentaStor User name and password.
 - For anonymous access, enter the appropriate anonymous User name and password.
5. Click **OK**.

Managing Access Control Lists

When you first share a folder, authenticated Users can log in to the folder with read-only access. You then create an ACL for each User or group that requires additional privileges for the folder.

Note: Use ACLs for authenticated Users only. ACLs are not required or used for anonymous User access.

ACLs are recursive unless you set the **Nonrecursive** property. You can manually create an ACL with different privileges for a child folder.

- ❖ *To set up an ACL for a User or group, using NMV in UNIX environment:*
 1. Click **Data Management > Data Sets**.
 2. In the **Folders** panel, click **Show**.
 3. In the **Summary Information: Folders** panel, click a shared folder.
 4. In the **Edit Folder: <fol_name> > Access Control Lists**, select one of the following:
 - (+) **Add Permissions for User**
 - (+) **Add Permissions for Group**
 5. In the **UNIX/LDAP Owner** or **UNIX/LDAP Group** field, type the name of the User or group.
 6. Select the permissions that you want to grant to the User or group.
 7. Click **Add New Group** or **Add New User**.

A new ACL entry for the User or group displays in the **Access Control List**, showing the User's or Group's permissions for that folder and all child folders (unless you set the non-recursive property).

- ❖ *To set up an ACL for a User or group, using NMC:*
 1. Type:

```
nmc:/$ setup folder <shared_fol_name> acl
```
 2. Select the entity type:
 - group
 - user
 - owner@
 - group@
 - everyone@
 3. Select the User or group.
 4. Select the permissions for the User or group.
 5. Press **Enter**.

Table 6-9: ACL Permissions

Permission	Description
add file	Add new file to a directory.
add subdirectory	Create a subdirectory in a directory.
delete	Delete a parent file.
delete child	Delete a file within a folder.
dir inherit	Inherit to all new child directories in a parent directory.
execute	Execute a file.
file inherit	Inherit to all new files in a directory.

Table 6-9: ACL Permissions

Permission	Description
inherit only	Applies the ACL privileges to all new files and sub-directories created within or copied to the selected folder. Use with <code>file_inherit</code> or <code>dir_inherit</code> . For new UNIX Users, Nexenta recommends not using <code>inherit only</code> .
list directory	List contents of a directory.
non-recursive	Grants the ACL privileges to the current folder only, and grants read-only access to child folders. You can create separate ACLs with custom privileges for child folders.
no propagate	Indicates properties should not be inherited by child folders.
read acl	Read the ACL file.
read attributes	Read the basic attributes of the file.
read data	Read the data of the file.
read xattr	Read the extended attributes of the file.
synchronize	Access a file locally at the server with synchronous reads and writes.
write acl	Write ACL or execute a <code>chmod</code> or <code>setfacl</code> .
write attributes	Change the times associated with a file or directory to an arbitrary value.
write data	Modify a file's data in the file offset range.
write owner	Change the owner or the ability to execute a <code>chown</code> or <code>chgrp</code> on a file.
write xattr	Ability to create extended attributes or write to the extended attribute directory.

Additional Folder Tasks

The following sections describe additional tasks for managing folders.

Unsharing a Folder

You can unshare a shared folder any time.

❖ *To unshare a folder, using NMV:*

1. Click **Data Management > Data Sets**.
2. In **Folders** panel, click **Show**.
3. In **Summary Information: Folders**, deselect the sharing option for the folder.
4. In the confirmation dialog, click **OK**.

❖ *To unshare a folder, using NMC:*

◆ Type one of the following:

```
nmc:/$ setup folder <vol_name/fol_name> unshare
```

```
nmc:/$ unshare folder <vol_name/fol_name>
```

Editing Folder Properties

You can edit folder's properties any time after you have created a folder.

❖ *To edit a folder's properties, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Folders** panel, click on Show.
3. In the **Summary Information: Folders**, click the folder's name.
4. In the **Read-Write Parameters** panel, edit the required properties.
5. Click **Save**.

❖ *To edit a folder's properties, using NMC:*

1. Display all of the available properties by typing:

```
nmc:/$ show folder <fol_name> property
```

2. Modify a particular property by typing the following command and following the prompts:

```
nmc:/$ setup folder <fol_name> property <propertyname>
```

Table 6-10: Folder Properties

Name (NMC)	Name (NMV)	Description
aclinherit	ACL Inheritance	Controls the inheritance, of the ACL settings, by new files and subdirectories from the parent directory. You can change these properties at any time. The options are: <ul style="list-style-type: none"> • discard — do not inherit the ACL entries • noallow — inherit ACL entries only with deny access type • restricted — inherit ACL entries, excluding write_owner, write_acl • passthrough — mode of newly created files are defined by inherited ACL entries • passthrough-x — files are created with a permission to execute, if this permission is defined in file creation mode and inherited ACL.

Table 6-10: Folder Properties

Name (NMC)	Name (NMV)	Description
<code>aclmode</code>	ACL Mode	Determines how to modify the ACL entry, when the system calls <code>chmod</code> . The options are: <ul style="list-style-type: none"> <code>discard</code> — remove all the ACL entries except the POSIX entries <code>owner@</code>, <code>group@</code> and <code>everyone@</code> <code>groupmask</code> — User permissions cannot be greater than group permissions. <code>passthrough</code> — ACL entries do not change when you change dataset permissions, except for the POSIX entries <code>owner@</code>, <code>group@</code>, <code>everyone@</code>
<code>atime</code>	Access Time	Determines whether to update the access time volume parameter every time you address files in this dataset. Setting this property to off can result in significant performance improvement, but can result in internal services confusion.
<code>canmount</code>	Can Mount	When set to on, you cannot mount the dataset with the <code>mount</code> command. Similar to setting <code>mountpoint</code> to <code>none</code> .
<code>checksum</code>	Checksum	Determines whether to use the proper checksum algorithm. If set to <code>on</code> , automatically uses the default checksum algorithm <code>fletcher4</code> .
<code>compression</code>	Compression	Enables the compression algorithm for a dataset.
<code>copies</code>	Number of copies	Determines the number of copies of user data for a dataset.
<code>dedup</code>	Deduplication	If set to on, enables the deletion of redundant data copies, using the storage capacity more effectively.
<code>devices</code>	Allow Devices	Controls whether device files in a file system can be opened.
<code>exec</code>	Allow Execution	Controls whether programs in a file system can be executed. Also, <code>mmap(2)</code> calls with <code>PROT_EXEC</code> are disallowed when set to off.
<code>groupquota</code>	N/A	Allocates the amount of space for a group can use in this folder.
<code>logbias</code>	Log Bias	Controls how ZFS optimizes synchronous requests for this dataset. If <code>logbias</code> is set to latency, ZFS uses the pool's separate log devices, if any, to handle the requests at low latency. If <code>logbias</code> is set to throughput, ZFS does not use the pool's separate log devices. Instead, ZFS optimizes synchronous operations for global pool throughput and efficient use of resources. The default value is latency.

Table 6-10: Folder Properties

Name (NMC)	Name (NMV)	Description
<code>mislabeled</code>	N/A	Provides a sensitivity label that determines if a dataset can be mounted in a Trusted Extensions zone. If the labeled dataset matches the labeled zone, the dataset can be mounted and accessed from the labeled zone. The default value is none. This property can only be modified when Trusted Extensions is enabled and only with the appropriate privilege.
<code>nbmand</code>	N/A	Controls whether the file system can be mounted with <code>nbmand</code> (Non-blocking mandatory) locks. This property is for SMB clients only. Changes to this property only take effect when the file system is unmounted and remounted.
<code>nms:description</code>	Description	Descriptive comment (optional).
<code>primarycache</code>	N/A	Controls what is cached in the primary cache (ARC). Possible values: <code>all</code> , <code>none</code> , and <code>metadata</code> . If set to <code>all</code> , it caches both user data and metadata. If set to <code>none</code> , it does not cache user data nor metadata. If set to <code>metadata</code> , it caches only metadata.
<code>quota</code>	Quota	Limits the amount of disk space a dataset and its descendents can consume. This property enforces a hard limit on the amount of disk space used, including all space consumed by descendents, such as file systems and snapshots. Setting a quota on a descendent of a dataset that already has a quota does not override the ancestor's quota, but rather imposes an additional limit. Quotas cannot be set on volumes, as the volume size property acts as an implicit quota.
<code>readonly</code>	Read-Only	Controls whether a dataset can be modified. When set to on, you cannot make modifications. Property abbreviation: <code>rdonly</code> .
<code>recordsize</code>	Record Size	Specifies a suggested block size for files in a file system. Property abbreviation: <code>resize</code> .
<code>refquota</code>	Referenced Quota	Sets the amount of disk space that a dataset can consume. This property enforces a hard limit on the amount of space used. This hard limit does not include disk space used by descendents, such as snapshots and <code>s</code> .

Table 6-10: Folder Properties

Name (NMC)	Name (NMV)	Description
refreservation	Referenced Reservation	Sets the minimum amount of disk space that is guaranteed to a dataset, not including descendents, such as snapshots and s. When the amount of disk space used is below this value, the dataset is treated as if it were taking up the amount of space specified by refreservation. The refreservation reservation is accounted for in the parent dataset's disk space used, and counts against the parent dataset's quotas and reservations. If refreservation is set, a snapshot is only allowed if enough free pool space is available outside of this reservation to accommodate the current number of referenced bytes in the dataset. Property abbreviation: refreserv.
reservation	Reservation	Sets the minimum amount of disk space guaranteed to a dataset and its descendents. When the amount of disk space used is below this value, the dataset is treated as if it were using the amount of space specified by its reservation. Reservations are accounted for in the parent dataset's disk space used, and count against the parent dataset's quotas and reservations. Property abbreviation: reserv.
secondarycache	N/A	Controls what is cached in the secondary cache (L2ARC). Possible values: all , none , and metadata . If set to all, it caches both user data and metadata. If set to none, it does not cache user data nor metadata. If set to metadata, it caches only metadata.
setuid	Allow Set-UID	Controls whether the setuid bit is honored in a file system.
snapdir	Snapshot Visibility	Controls whether the .zfs directory is hidden or visible in the root of the file system

Table 6-10: Folder Properties

Name (NMC)	Name (NMV)	Description
<code>sync</code>	Sync	Determines the synchronous behavior of a file system's transactions. Possible values: <ul style="list-style-type: none"> • Standard: It writes synchronous file system transactions, such as <code>fsync</code>, <code>O_DSYNC</code>, <code>O_SYNC</code>, and so on, to the intent log. (default value) • Always: Ensures that it writes and flushes every file system transaction to stable storage, by a returning system call. This value carries a significant performance penalty. • Disabled: Disables synchronous requests. It only commits file system transactions to stable storage on the next transaction group commit, which might be after many seconds. This value provides the best performance, with no risk of corrupting the pool. However, this value is very dangerous because ZFS is ignoring the synchronous transaction demands of applications, such as databases or NFS operations. Setting this value on the currently active root or <code>/var</code> file system might result in out-of-spec behavior, application data loss, or increased vulnerability to replay attacks. Only use this value if you fully understand the risks.
<code>userquota</code>		Specifies storage space quota for a particular user.
<code>version</code>	N/A	Identifies the current version of the volume
<code>vscan</code>	N/A	Controls whether regular files should be scanned for viruses when a file is opened and closed. In addition to enabling this property, a virus scanning service must also be enabled for virus scanning to occur if you have third-party virus scanning software. The default value is off.
<code>xattr</code>	Extended Attributes	Indicates whether extended attributes are enabled (on) or disabled (off) for this file system.

Deleting a Folder

You can delete/destroy a NexentaStor folder anytime after its creation. When you delete a folder, the folder is automatically unshared and unmounted.

❖ *To delete a folder, using NMV:*

1. Click **Data Management > Data Sets**.
2. In **Folders** panel, click **Show**.
3. In **Summary Information: Folders** window, select the folders for deletion.
4. Click **Delete Selected**.
5. Alternatively, click the cross icon under **Delete** to delete a single folder.

6. In the confirmation dialog, click **OK**.

Note: If you try to delete a folder in which a process is running, the operation fails. When this happens, you must destroy the folder to delete it.

Caution: Forcing the operation can result in unpredictable system failures and data inconsistency. Use force with caution.

You can also delete the folders after filtering them.

❖ *To delete a folder, using NMC:*

◆ Type one of the following:

```
nmc:/$ destroy folder <fol_name>
nmc:/$ setup folder <fol_name> destroy
System response:
Destroy folder 'vol1/foll' ? (y/n) y.
Folder 'vol1/foll' destroyed.
```

See Also:

- [Filtering Folders](#)

Indexing a Folder

To include a folder in search results, you must first index it.

Runner Indexer executes according to a schedule and collects the information about indexed folders and their snapshots. To make the search available, you can index the commonly used folders.

The NexentaStor Indexer uses phrase and a proximity searching method. Indexer supports the indexing of the following file formats:

- AbiWord
- CSV
- Debian packages
- DVI
- HTML
- Microsoft Excel
- Microsoft Powerpoint
- Microsoft Word
- Microsoft Works
- MySQL data
- OpenDocument
- OpenOffice
- PDF
- Perl POD documentation
- PHP
- PostScript
- RDBMS data
- RPM packages
- RTF
- StarOffice
- SVG
- text
- Word Perfect

❖ *To index a folder, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Folders** panel, click **Show**.
3. In the **Folder Summary View** window, select the folders to **Index**.
4. In the confirmation dialog, click **OK**.

❖ *To index a folder, using NMC:*

1. Create an indexer for the folder by typing one of the following:

```
nmc:/$ create indexer <fol_name>
nmc:/$ setup indexer create
```

2. Run the indexer by typing one of the following:

```
nmc:/$ run indexer <fol_name>
nmc:/$ setup indexer <fol_name> run
```

If you have created more than one indexer, you can run all indexers at once.

❖ *To run all indexers simultaneously:*

- ◆ Type one of the following:

```
nmc:/$ run indexer
nmc:/$ setup indexer run
```

Searching a Folder

You can search a folder only after you index it.

After setting up an index on a folder, the Runner Indexer must run at least once to index all of the folders and make them display in search results.

❖ *To search for a folder:*

1. Click **Data Management > Data Sets**.
2. In the **Folders** panel, click **Search**.
3. In the **Search** window, type in the corresponding fields to start the search.

See Also:

- [Managing Folders](#)
- [Indexing a Folder](#)

Filtering Folders

You can filter the folders to narrow the parameters and more efficiently perform the required operations.

❖ *To filter the folders:*

1. Click **Data Management > Data Sets**.
2. In the **Folders** panel, click **Show**.

3. In the **Summary Information: Folders** panel, type the required filtering parameter and click **Filter**.

Tip:

After filtering the folders a new button, **Delete filtered**, displays next to the **Filter** button. Use it to delete all of the folders with the parameters specified in the **Filter** field.

Sharing Folders Using CIFS

This section includes the following topics:

- [About CIFS](#)
- [Choosing the Mode](#)
- [Workgroup Mode Authentication](#)

About CIFS

This document explains how to use the Common Internet File System (CIFS) capabilities to share NexentaStor folders for:

- Anonymous access
- Authenticated access in:
 - Workgroup mode
 - Domain mode

CIFS service operational mode is system-wide, and it is either a workgroup or a domain. Meaning, NexentaStor cannot provide some CIFS shares to workgroup Users and, simultaneously, other shares to Users joined through Active Directory.

By default, NexentaStor operates in workgroup mode. The default pre-configured workgroup name is: WORKGROUP.

See Also:

- <http://en.wikipedia.org/wiki/CIFS>

Choosing the Mode

The system administrator decides which mode best matches the company's network configuration needs.

Workgroup Mode

Typically, small companies or home networks use workgroups. They are a group of loosely connected computers, through a peer-to-peer network. This means that each computer is sustainable on its own. It has its own User list, its own access control and its own resources. In order for a User to access resources on another workgroup computer, that exact User must be created on the other computer. This method is simple to design and implement, but if your network grows, it becomes difficult to manage.

Example:

A User would need an account on all of the computers it needs to access. Any account changes, (i.e. password changing) must be applied to all of the computers in the workgroup. It is not applicable for a network with more than 50 computers.

Workgroup mode:

- Applies in small networks (less than 10 computers)
- Easy to setup and does not require any additional knowledge
- Requires setting up accounts and passwords on each and every computer

Domain Mode

A domain is a trusted group of computers that share security, access control and have data passed down from a centralized domain controller server or servers. Domain mode requires additional arrangements on the Windows side. It requires a configured Domain Controller with DNS (Windows Server 2003/2008) which manages all of the aspects of granting User permission to a login. Domain mode is commonly used in large networks and provides advanced centralized management and security, but it is more complex in design and implementation than a workgroup mode setup.

Domain mode:

- Single location for all User accounts, groups and computers, passwords are the same for all computers.
- Requires configured Domain Controller (or two: primary and backup) with Active Directory and DNS server.
- More difficult to set up and requires additional knowledge.

If you use an appliance's CIFS for anonymous access, authenticated (workgroup) or in domain mode, you must configure a CIFS Server on NexentaStor.

Workgroup Mode Authentication

The CIFS server acquires authentication through local authorization when in Workgroup mode.

- ❖ *To join a local workgroup, using NMV:*
 1. Click **Data Management > Shares**.
 2. In the CIFS Server panel, click **Join Workgroup**.
 3. Optionally, modify the name of the workgroup.
 4. Click **Save**.

❖ *To join a workgroup, using NMC:*

1. Type:

```
nmc:/$ setup network service cifs-server join__workgroup
```

System response:

```
Workgroup Name : WORKGROUP
```

2. Modify the name of the workgroup if needed.

See Also:

- [About Using Windows Active Directory](#)

This page intentionally left blank

Using Windows Active Directory

This section includes the following topics:

- [About Using Windows Active Directory](#)
- [Prerequisites](#)
- [Verifying DNS Settings](#)
- [Adding NexentaStor to Active Directory](#)
- [Setting Permissions](#)
- [Domain Mode Authentication](#)
- [About ID Mapping](#)
- [Mapping Windows Users to NexentaStor Users](#)
- [Adding Members to Groups](#)
- [Removing Members from Groups](#)
- [Modifying the Access Control List.](#)
- [Troubleshooting](#)

About Using Windows Active Directory

You can integrate NexentaStor with Windows Active Directory to take advantage of Active Directory security policies on shared files and directories. You can integrate it in a simple network with a few Active Directory servers, or a large network with multiple Active Directory servers.

You may optionally create ID mappings for AD users and groups. These optional ID mappings are helpful for NFS clients in a mixed NFS and SMB environment. In an environment with only SMB clients, ID mappings are not recommended.

- [About ID Mapping](#)
- [Domain Mode Authentication](#)
- [Prerequisites](#)

Prerequisites

Before starting the Active Directory integration, make sure the environment meets the following prerequisites:

- Windows Server 2003 or later version.
- Active Directory configured on Windows Server.
- DNS server installed and working in the Active Directory environment.
- NexentaStor and the Domain Controller use the same NTP Server and the time is in sync on both systems.
- If the NexentaStor network interface is configured as DHCP, the DHCP server name-server list contains the IP address of the DNS server that is used for the Active Directory domain.

See Also:

- [Sharing Folders Using CIFS](#)

Verifying DNS Settings

Before you integrate NexentaStor with Active Directory, verify that the DNS settings on the NexentaStor system are correct. Change any settings as needed.

To complete this procedure, you need the following information:

- The name of the Active Directory domain
- The IP address of a DNS server within the Active Directory environment

❖ *To verify the DNS settings, using NMC:*

1. Using an ssh client, log in to NMC as root.

2. Type:

```
nmc:/$ show network service cifs-server config
```

3. Select `resolv.conf`.

Example:

```
nmc:/$ show network service cifs-server config
```

System response:

```
cifs-server configuration file : resolv.conf
search nexenta.com
nameserver 10.3.250.2
```

4. If you need to make any changes, edit the `resolv.conf` file using the following command:

```
nmc:/$ setup network service cifs-server edit-settings resolv.conf
```

Note: Note the DNS domain can be different from the AD domain

See Also:

- [Prerequisites](#)
- [Adding NexentaStor to Active Directory](#)

Adding NexentaStor to Active Directory

To complete this procedure, you need the following information:

- The nameserver and domain values from the `resolv.conf` file (see [Verifying DNS Settings](#))
- Either the Domain Administrator, or a User with administrative privileges, can add and create new computer objects in the domain.

❖ *To add NexentaStor to Active Directory, using NMC:*

1. Click **Data Management > Shares**.
2. In the **CIFS Server** panel, click **Join AD/DNS Server**.
3. In the **Manage CIFS Server Settings** window, type the following information:
 - **AD Domain**
AD domain name need not match the DNS domain
 - **AD Join User**
Name of a User who is part of the Active Directory Domain Administrators group, with privileges to create new computer objects
 - **AD Join Password**
Active Directory password for the AD Join User
4. Click **Save**.
A success message displays at the top of the screen.
5. Using NMC, verify that the CIFS server can see NexentaStor by typing the following command:

```
nmc:/$ show network service cifs-server
```

See Also:

- [Prerequisites](#)
- [Setting Permissions](#)
- [Troubleshooting](#)

Viewing Service Logs

If NexentaStor failed to join the Active Directory, view the service logs to troubleshoot.

❖ *To view the Service Logs, using NMV:*

1. Click **Data Management > Shares**.
2. In the **CIFS Server** panel, click **View Log > View Service Logs**

❖ *To view the service logs, using NMC:*

1. Type:

```
nmc:/$ show log syslog
```

Setting Permissions

Follow this procedure on a Windows client in the Active Directory domain if you are an user who has Administrative privileges.

❖ *To set permissions for the Active Directory users and groups, using Windows:*

1. Go to **Start > Run**.
2. In the **Run** dialog, type the UNC path for the NexentaStor appliance. An example is `\\myhost`.
The shared folders display.
3. Right-click on a shared folder and select **Properties**.
4. In the **Properties** dialog, click the **Security** tab.
5. Select a group or User and click **Allow** or **Deny** to set the permissions for that group or user.
6. Repeat the previous steps until appropriate permissions are set for all Users and groups of all shared folders.

Obtaining Administrative Privileges

In domain mode any domain user who is a member of the Administrators group has administrative privileges. The following accounts have administrative privileges by default:

- `local\root`
- `local\admin`
- Any member of the local Administrators group.

To give other groups administrative privileges, edit the membership of the local Administrators group. For example, if you have a domain group named **nexentastor-admins**, you add that to the local Administrators group as described in section [Adding Members to Groups](#)

❖ *To obtain the administrative privileges for a user, using Windows Active Directory:*

1. Remove any mappings to the server.
2. Select **Connect as different user**.
3. Choose an account with administrative privileges.

See Also:

- [Verifying DNS Settings](#)
- [Modifying the Access Control List](#).

Domain Mode Authentication

The domain mode method of user authentication involves integrating NexentaStor with Active Directory so that NexentaStor can use the centralized Active Directory authentication and authorization information.

Before you start, make sure you have the following:

- Windows 2003 Server with Active Directory configured, or Windows 2008 Server SP2, or a later version with Active Directory configured.
- DNS server installed and working as part of the Active Directory environment, including zone configuration, proper delegations in parent DNS zones, and DNS domain controller locator records (SRV records).

See Also:

- [Joining Active Directory](#)
- [Initial Troubleshooting Steps](#)
- [General Troubleshooting](#)

Joining Active Directory

There are two different scenarios of integrating NexentaStor with Windows Active Directory:

Configuring Windows

Windows contains two types of access.

- **Full Control**

If the NexentaStor appliance is already registered with Active Directory, you can use any valid User account. The account must have Full Control over this particular appliance, but does not need administrative privileges.

- **Standard Control**

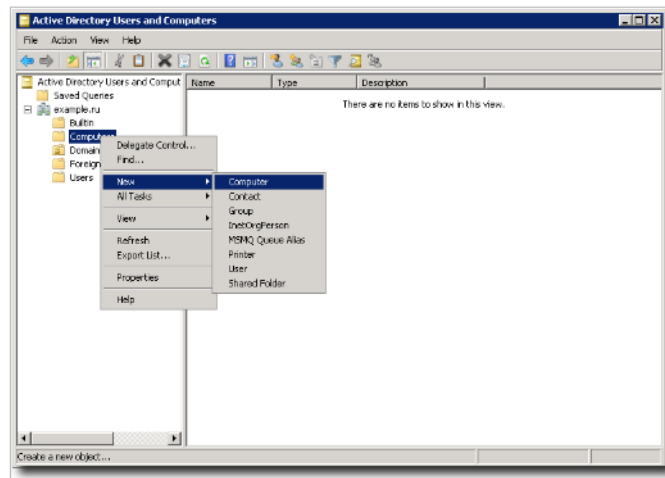
If there is not already a NexentaStor appliance, registered with Active Directory, the Windows Administrator must create a NexentaStor computer object before proceeding.

The following assumes that NexentaStor appliance is not yet present in the Active Directory database. In this scenario, the Windows Administrator must create a corresponding computer object.

Note: NexentaStor DNS entry is not automatically added to the Active Directory DNS server. After NexentaStor joins the Active Directory, you must add the NexentaStor IP address and hostname record to DNS server.

❖ *To create a NexentaStor computer object for Active Directory, using Windows:*

1. In Windows Server, start **Microsoft Management Console**.
2. Right click on Computers, and select **New**.



3. For computer name, specify the NexentaStor appliance by hostname.
4. Right-click on the computer you just added and select **Properties**.
5. Click **Security** tab.
6. Type in the user (or group) name.
7. Click **Check Names**.

Make sure to provide the newly added computer Users with Full Control over this computer.

Note: You can skip using Microsoft Management Console and performing Steps 1 through 4 (above) in either one of the following two cases:

- (1) Account with administrative privileges is used to perform a join operation.
- (2) A record of the computer object representing the appliance already exists.

The rest of this section assumes that either (1) or (2) above, or both, are true.

See Also:

- [General Troubleshooting](#)

Joining a NexentaStor Appliance to an Active Directory Domain

❖ *To join a NexentaStor appliance to an Active Directory domain, using NMV:*

1. Go to **Settings > Misc Services**.
2. Click **Join AD/DNS Server**.

❖ *To join a NexentaStor appliance to an Active Directory domain, using NMC:*

◆ **Type:**

```
nmc:/$ setup network service cifs-server join_ads
```

Example:

```
DNS Server IP address, port : 172.16.44.182
AD Domain Name             : example.ru
AD Join User                : Administrator
AD Join Password           : xxxxxxxxxx
```

A successful join is persistent across reboots.

Moving NexentaStor Appliance from one AD Domain to Another AD Domain

You may need to switch a NexentaStor appliance from one AD domain to another AD domain if the name of the current domain changed or a new domain has been configured in your network environment. To join a new domain, first switch NexentaStor to a workgroup mode and then join the new domain.

❖ *To become a member of a workgroup, using NMV*

1. Go to **Data Management > Shares**.
2. In the **CIFS Server** panel, click **Join Workgroup**.
3. In the **Manage CIFS Server Settings** window, change the workgroup name, if needed.
4. Click **Save**.
5. Go to **Data Management > Shares**.
6. In the **CIFS Server** panel, click **Join AD/DNS server**.
7. In the **Manage CIFS Server Settings** window, type the following information:

- **AD Domain**

AD domain name need not match the DNS domain

- **AD Join User**

Name of a User who is part of the Active Directory Domain group, with privileges to create new computer objects

- **AD Join Password**

Active Directory password for the AD Join User

8. Click **Save**.

A success message displays at the top of the screen.

See Also:

- [Troubleshooting](#)
- [Sharing a Folder Using CIFS](#)

About ID Mapping

While enabling Active Directory users to access NexentaStor NFS and SMB folders, you may establish the user name mapping between Windows Users and groups and their counterparts in UNIX through the appliance's `idmap` facility. For SMB no ID mapping needed. These optional ID mappings are helpful for a mixture of NFS and SMB clients. The `idmap` mappings persist across reboots.

The identity mapping utility supports user to user, group to group, group to user, and user to group mappings.

If you create a group to group mapping, you must first create and map all users from a Windows group on the NexentaStor appliance. Then, you must create a NexentaStor group, include all the users to that group, and then create a group to group mapping.

To facilitate the process of mapping Active Directory users to NexentaStor users, you can map a Windows group to a single NexentaStor user.

Mapping Windows Users to NexentaStor Users

Note:

Using `idmap` rules with wildcard match patterns affect performance. You can use wildcards for NFS and CIFS environments where there is a requirement to have normal-looking user IDs exposed through NFS. In a CIFS-only environment, you should never need to add `idmap` rules.

NexentaStor uses the `idmap` utility to associate Windows user and group security identifiers (SID) with UNIX user and group identifiers (UIDs and GIDs). You must map the Windows Active Directory users to the NexentaStor users to enable them to access the NexentaStor CIFS folders with the configured privileges.

❖ *To map Windows users/groups onto UNIX Users/groups, using NMV:*

1. Click **Data Management > Shares**.
2. In the CIFS Server panel, click **Identity Mapping**.
 - If you map a user:
 - 1) Select **winuser**.
 - 2) Type the account name from Windows Active Directory.
 - 3) Depending on the type of mapping, select **unixuser** or **unixgroup**.
 - 4) Type the NexentaStor user or group.

The user or group must exist on the NexentaStor appliance.

5) Click **Add Rule**.

- If you map a group:

1) Select wingroup.

2) Type the name of group from Windows Active Directory.

3) Depending on the type of mapping, select **unixuser** or **unixgroup**.

4) Type the NexentaStor user or group.

The user or group must exist on the NexentaStor appliance.

5) Click **Add Rule**.

❖ *To map Windows Users/groups onto UNIX Users/groups, using NMC:*

1. Type:

```
nmc:/$ setup network service cifs-server idmap
```

2. Specify the Windows User name by using one of the following formats:

```
winuser:username@domain-name
```

```
winuser:domain-name\username
```

3. Specify a Unix User name in the following format: `unixuser:username`

Note: Windows User names are not case insensitive, while Illumos User names are case sensitive.

Example:

Map all users in the domain mydomain.com:

```
winuser:*@mydomain.com==unixuser:*
```

Map Unix user joe to Windows user Joe in the domain mydomain.com:

```
winuser:Joe@mydomain.com==unixuser:joe
```

Note: Using idmap rules with wildcard match patterns affect performance. You can use wildcards for NFS and CIFS environments where there is a requirement to have normal-looking user IDs exposed through NFS. In a CIFS-only environment, you should never need to add idmap rules.

There are popular Windows user and group names that are supported by idmap:

- Administrator
- Guest
- KRBTGT
- Domain Admins
- Domain Users

- Domain Guest
- Domain Computers
- Domain Controllers

When you add `idmap` rules, these popular names are expanded to canonical form. This means that either the default domain name is added (for names that are not popular) or an appropriate built-in domain name is added. Depending on the particular name, the domain name might be null, built-in, or the local host name.

Example:

If you map `wingroup Administrators` to `unixgroup sysadmin`, it automatically maps with the `@BUILTIN` virtual domain:

```
nmc:/$ setup network service cifs-server idmap
nmc:/$ Mappings Rules : wingroup:Administrators==unixgroup:sysadmin
nmc:/$ show network service cifs-server idmap
nmc:/$ add wingroup:Administrators@BUILTIN unixgroup:sysadmin
```

See Also:

- [About ID Mapping](#)
- [Modifying the Access Control List.](#)

Default CIFS Accounts

NexentaStor contains default Unix and Window accounts that you MAY map to each other for use with NFS and CIFS environment. ID mapping is not recommended for CIFS only environment.

Table 8-1: Mapping Default Accounts

Account Type	Windows Account	Unix Account	Mapping Command	Initial Members
Local CIFS Administrator	Administrator	admin	<code>idmap add winuser:Administrator * unixuser:admin</code>	Administrators root user admin user
Local CIFS Guest	Guest	smb	<code>idmap add winuser:Guest* unixuser:smb</code>	Guest
Local CIFS Users group	Guest	smb	<code>idmap add -d unixuser:smb wingroup:Guests</code>	
Local Administrators group	Windows Domain Administrators	admin	<code>idmap add -d unixuser:admin wingroup:Administrators</code>	

Table 8-1: Mapping Default Accounts

Account Type	Windows Account	Unix Account	Mapping Command	Initial Members
Local Users group	Administrators	root	idmap add -d unixuser:root wingroup:Administrato rs	
Local Guests	Windows Users	staff	idmap add wingroup:Users unixgroup:staff	Domain Guests

Adding Members to Groups

You can add members to Windows or Unix domains.

- ❖ *To add a user to a Nexenta domain, using NMC:*

- ◆ Type:

```
nmc:/$ smbadm add-member -m "administrator@<NexentaStor domain>"
"administrators"
```

- ❖ *To add a user to an AD domain, using NMC:*

- ◆ Type:

```
nmc:/$ smbadm add-member -m "<AD domain>\administrator" "administrators"
```

Example:

```
nmc:/$ smbadm add-member -m "MyDomain/IT Staff" "backup operators"
```

```
nmc:/$ smbadm add-member -m "MyDomain/IT Staff" "power users"
```

See Also:

- [Removing Members from Groups](#)

Removing Members from Groups

You can delete members to Windows or Unix domains.

- ❖ *To remove a user from an AD domain, using NMC:*

- ◆ Type:

```
nmc:/$ smbadm remove-member -m <AD domain>\administrator administrators
```

Table 8-2: Mapping Commands

Command	Description
get_idmap	Show the list of idmaps.
do_idmap	Clear the idmaps.
set_idmap in	Set the idmap.
smbadm_join_ad	Add new mappings to an Active Directory.

See Also:

- [Adding Members to Groups](#)
- [Modifying the Access Control List.](#)

Modifying the Access Control List.

You can modify access permissions for Windows Active Directory users in NMV, as well as in Windows Active Directory. You can also add new Windows users and groups to the CIFS folder Access Control List (ACL). The Windows users and groups must be mapped to UNIX users and groups on the NexentaStor appliance. Access control settings in NexentaStor are "unified" so when you modify the ACL in NexentaStor, changes made are visible to Windows clients, and vice versa.

❖ *To modify the Access Control List, using NMV:*

1. Click **Data Management > Shares**.
2. Click on the CIFS share.
3. Select the action:
 - **Add Permissions for User.**
 - 1) Type the UNIX user name.
The UNIX user must exist on the NexentaStor appliance.
 - 2) Modify the access permissions for the selected user.
 - 3) Click **Add New User**.
 - **Add Permissions for Group.**
 - 1) Type the UNIX group name.
The UNIX group must exist on the NexentaStor appliance.
 - 2) Modify the access permissions for the selected group.
 - 3) Click **Add New Group**.

- **Reset ACL to Defaults.**

This option resets the ACL list to default. The default ACL list includes the following users:

- owner
- group
- everyone

See Also:

- [Managing Folders](#)
- [About ID Mapping](#)

Modifying ACL in Windows

When you modify the access control settings for Windows clients, you can use a Windows client to edit the settings as described in [Setting Permissions](#).

❖ *Modifying ACLs in Windows:*

1. Log in to Windows as a user with administrative privileges.
2. Right click on the folder you want to modify.
3. Click the **Security** tab.
4. Change settings.
5. Click **Save**.

Troubleshooting

The following sections include various troubleshooting strategies and tips.

Initial Troubleshooting Steps

NexentaStor displays a message to indicate a successful join, or a failure to join an Active Directory.

If you fail to join an AD, you can troubleshoot the problem.

❖ *To troubleshoot Active Directory, using NMV:*

- ◆ Click **Settings > Misc. Services > View Log**

❖ *To troubleshoot Active Directory, using NMC:*

1. Type:

```
nmc:/$ show network service cifs-server log
```

System response:

- network-smb-server:default.log
- messages

2. Select messages .

The following displays an example of the messages log:

```
Dec  3 03:11:34 nexentastor idmap[355]: [ID 523480 daemon.notice] AD
loOkup of winname Administrator@John-PC failed, error code -9961
Dec  3 03:12:29 nexentastor last message repeated 7 times
Dec  3 03:12:54 nexentastor smbd[374]: [ID 812811 daemon.notice]
logon[John- PC\alice]: WRONG_PASSWORD
```

General Troubleshooting

The following troubleshooting tips are common for all versions of Windows Servers:

1. Make sure the time is in sync using the same NTP Server for both Domain Controller and NexentaStor.
2. Verify that DNS is properly configured.

Verify that DNS is configured properly by ensuring that both `domain` and `search` parameters point to the Active Directory domain name. Parameter `nameserver` must have the IP address of a DNS server within the Active Directory environment.

❖ *To check the configuration, using NMC:*

◆ Type:

```
nmc:/$ show network service cifs-server config
```

If you need to apply any corrections, you must make the changes in VI editor.

❖ *To edit the file in vi editor:*

◆ Type:

```
nmc:/$ setup network service cifs-server edit-settings resolv.conf
```

Note:

If you configured the network interface as DHCP, the DHCP server's `name-servers` list must contain the DNS server which is used for the domain. Otherwise, the list is updated after a reboot and the AD connection is lost.

Verifying SRV Records of Domain Controller

When NexentaStor requests access to AD, new DC locator uses the netlogon/winlogon service to locate an AD server by querying the DNS Server to find the IP Address of the domain controller that is used by the SMB service, as well as the ID mapping service. Locators find the AD servers by using DNS names (for IP or DNS-compatible computers). SRV records play an important role for domain controllers in the Active Directory domain. You will not know the location of a domain controller without the SRV records.

❖ *To verify records of Domain Controller, using NMC:*

1. Log in to bash:

```
nmc:/$ option expert_mode =1
nmc:/$ !bash
```

2. Type the following command (note the space before `_ldap`):

```
# dig @<nameserver> _ldap._tcp.dc._msdcs.<domain> SRV +short
```

Example:

```
# dig @10.3.250.2 _ldap._tcp.dc._msdcs.eng.nexenta.corp SRV +short
```

3. Verify that the response is the domain controller for the Active Directory domain:

Example:

```
# dig @10.3.250.2 _ldap._tcp.dc._msdcs.eng.nexenta.corp SRV +short
  0 100 389 eng.nexenta.com.
```

4. Switch to NMC:

```
# exit
```

5. Verify the NexentaStor has Joined the domain:

```
nmc:/$ show network service cifs-server
```

Active Directory Integration

If problems occur during Active Directory integration, see the following issues and resolutions.

Table 8-3: Active Directory Integration

Problem	Suggested Solution
After rebooting NexentaStor, Active Directory is no longer connected.	If the NexentaStor network interface is configured as DHCP, make sure the DHCP server nameserver list contains the IP address of the DNS server that is used for the Active Directory domain. This list is updated when NexentaStor reboots.

Unusual User/Group IDs in ACL

After NexentaStor joins Active Directory and you modify the ACL for the NexentaStor CIFS folder, you may see unusual user and group entities in the CIFS share ACL list:

Entity	Allow	Deny	Delete
owner@	list_directory, read_data, add_file, write_data, add_subdirectory, append_data, read_xattr, write_xattr, execute, read_attributes, write_attributes, read_acl, write_acl, write_owner, synchronize		-
group@	list_directory, read_data, read_xattr, execute, read_attributes, read_acl, synchronize		-
everyone@	list_directory, read_data, add_file, write_data, add_subdirectory, append_data, read_xattr, write_xattr, execute, read_attributes, write_attributes, delete, read_acl, synchronize, file_inherit, dir_inherit		-
group:2147483654	list_directory, read_data, add_file, write_data, add_subdirectory, append_data, read_xattr, write_xattr, execute, delete_child, read_attributes, write_attributes, delete, read_acl, write_acl, write_owner, synchronize, file_inherit, dir_inherit		✗
user:2147483652	list_directory, read_data, add_file, write_data, add_subdirectory, append_data, read_xattr, write_xattr, execute, read_attributes, write_attributes, delete, read_acl, synchronize, file_inherit, dir_inherit		✗
user:smb	list_directory, read_data, add_file, write_data, add_subdirectory, append_data, read_xattr, write_xattr, execute, delete_child, write_attributes, write_acl, write_owner, synchronize		✗
user:2147483650	list_directory, read_data, read_xattr, execute, read_attributes, read_acl, synchronize, file_inherit, dir_inherit		✗
user:2147483651	list_directory, read_data, add_file, write_data, add_subdirectory, append_data, read_xattr, write_xattr, execute, read_attributes, write_attributes, read_acl, synchronize, file_inherit, dir_inherit		✗

You may see these entities if you did not create ID mappings for these users. After you create ID mappings for Windows users, NexentaStor automatically replaces the unexpected numeric User/Group IDs and replaces them with proper names.

Note:

To view User/Group Name instead of the unusual IDs, configure the Operating System name services using `/etc/nsswitch.conf` to resolve information about the users and the group.

This page intentionally left blank

Managing Users

This section includes the following topics:

- [About Users and User Permissions](#)
- [About Default NexentaStor Users](#)
- [NMV User Permissions](#)
- [About Local Appliance User Groups](#)
- [About LDAP Users](#)

About Users and User Permissions

You can create as many NexentaStor user accounts as needed. Each user has an individual login and password. Each user is also automatically assigned a unique user ID (visible in NMC only).

By default, each user has permission to view all data. You can assign additional permissions to users, such as permission to search folders, add runners, import datasets, and so on.

You can also create a user that only has access to using NexentaStor through NMV. This User cannot use NexentaStor through a Unix shell or NMC.

About Default NexentaStor Users

NexentaStor provides the following predefined Users.

Table 9-1: NexentaStor Predefined Users

User	Description
root	Super-User account that allows you to perform any action in NMC. Only use this account for NMC administration. You cannot log on to NMV as root.
admin	Administrative account for NMV management. Admins can view all NMV pages and perform all NMV actions. Caution: Logging on to NMC as an admin opens a UNIX shell. Nexenta recommends that only advanced Users use the UNIX shell.
guest	User account with no administrative privileges. By default, guests can only view non-confidential data.

NMV User Permissions

In NMV, each User can have an individual set of access permissions. In NMC, each User is assigned to a group, and inherits the access privileges granted to the group.

You can change these permissions in NMV in the **Settings > Users** screen.

Note: In addition to granting User permissions, you can also configure User authentication. User authentication allows Users to access and perform actions on shared folders. For information about User authentication, see [About Users and User Permissions](#).

Table 9-2: NMV Permissions

Permission	Administration Area	Location in NMV
Unix User	If not selected, User cannot use NMC. User only has NMV account.	
Unix group	Default group User belongs to.	
Home Folder	Home folder path name.	
Can_view	View-only access to most NMV pages, settings, logs, and other data	Most screens
Can_admin_groups	Groups of appliances	Status > General > Appliance groups
Can_admin_services	Auto-services	Data Management > Autoservices
Can_admin_runners	Runners	Data Management > Runners
Can_admin_datasets	Datasets	Data Management > Data sets
Can_admin_shares	Full administrative access to folders shared using CIFS, NFS, and FTP	Data Management > Shares
Can_admin_users	Users and Access Control lists	Settings > Users
Can_admin_disks	Disks	Settings > Disks
Can_admin_network	Network settings and miscellaneous services	Settings > Network Settings > Misc Services
Can_admin_appliance	License key and various properties Power off and reboot	Settings > Appliance Analytics > Settings
Can_admin_prefs	Various preferences	Settings > Preferences Analytics > Settings

Table 9-2: NMV Permissions

Permission	Administration Area	Location in NMV
Can_search_folders	Searching folders and snapshots for content	Data Management > Data sets
Can_use_restapi	REST-API	

Creating Users and Setting Permissions

- ❖ *To add and set permissions for a User, using NMV:*
 1. Click **Settings > Users**.
 2. In the **Users** panel, click **New User**.
 3. In the **Create New Appliance User** window, type the appropriate information for this User.
 4. (Optional) Specify a home folder for the User's files.
 5. Click **Save**.
 6. In the **NMV GUI Permissions** window, select the appropriate NMV permissions for this User. For a list of permissions, see [NMV User Permissions](#).
 7. Click **Create New**.
- ❖ *To add and set permissions for the User, using NMC:*
 1. Type one of the following:

```
nmc:/$ setup appliance user create
nmc:/$ create appliance user
```
 2. (Optional) Specify a home folder for the User's files.
 3. Select a local appliance group for the User.

The User inherits the permissions granted to this group.
 4. Type and confirm the password for the User.

NexentaStor creates the User account with the specified password.

Setting the Quota for a User Folder

You can set and change the size limit on a folder created for a User.

Note:	This action is only available using NMC.
Warning:	If you change the quota to <code>none</code> , NexentaStor deletes all of the existing data in the User's folder.

❖ *To set a quota for a User folder, using NMC:*

1. Type:

```
nmc:/$ setup folder <vol_name/fol_name> property userquota
```

2. Type the User name.

3. Type the quota size.

You can display all User folders and the quota for each.

❖ *To list all User folders and their quota, using NMC:*

◆ Type:

```
nmc:/$ show folder <vol_name/fol_name> property quota
```

Managing Users

You can add, delete, and modify NexentaStor Users using NMV in the **Settings > Users** screen.

Table 9-3: NMC User Management Commands

Command	Action
create appliance user setup appliance user create	Create a new User
show appliance user	List all Users. For each User, display the User ID, type, group, home directory, and description.
setup appliance user <user_name> property uidNumber	Change a User ID number. NexentaStor automatically syncs User IDs in the /etc/passwd and /var/smb/smbpasswd files. If you change a User ID, the User cannot access the files in their directory until you change the file ownership to the new User ID.
create appliance user setup appliance user create	Create a new User

About Local Appliance User Groups

Local appliance User groups (also called User groups or local appliance groups) allow you to easily manage User permissions. You create a group and assign permissions to the group. Later, when you create Users, you specify a group for each User. The User inherits the permissions of the group.

You can create as many User groups as needed. Each User group can have as many Users as needed.

NexentaStor provides the following pre-defined groups:

- The `root` group contains the `root` User by default, and provides NMC super User permissions. All Users that you assign to this group have super User permissions.
- The `staff` group contains the `admin` User by default, and provides NMV administrative permissions. All Users you assign to this group have permission to perform all NMV tasks.
- The `other` group does not contain any Users by default. All Users you assign to this group have read-only permissions and cannot make any changes using NMV or NMC.

Creating Local User Groups

You can create local User groups. NexentaStor creates a unique ID number for each group.

- ❖ *To create a new appliance group, using NMV:*
 1. Click **Settings > Users**.
 2. In the **Groups** panel, click **New Group**.
 3. In the **Create New Appliance User Group** window, in the Group Name field, type the name of a new group.
 4. (Optional) In the **Available Users** list, choose the Users to add to this group and click **Add**.
 5. Click **Create New**.
- ❖ *To create a new appliance group, using NMC:*
 1. Type:


```
nmc:/$ setup appliance usergroup create
nmc:/$ create appliance usergroup
```
 2. Type the group name.
 3. Select one or more members to add to the group.

Setting the Quota for a Group

You can select a folder for use by a User group and set a quota/size limit on the folder.

Note: This action is only available using NMC.

Warning: If you change the quota to `none`, NexentaStor deletes all of the existing data in the User's folder.

- ❖ *To set a quota on a group folder, using NMC:*
 1. Type:


```
nmc:/$ setup folder <vol_name/fol_name> property groupquota
```
 2. Type the group name.
 3. Type the quota size.

You can list all group folders and the quota for each.

❖ *To list all group folders and the quota for each folder:*

◆ Type:

```
nmc:/$ show folder <vol_name/fol_name> property groupquota
```

Managing Local User Groups

You can add, delete, and modify NexentaStor groups in NMV on the **Settings > Users** screen.

You can manage NexentaStor groups in NMC with the following commands.

Table 9-4: NMC Group Management Commands

Command	Action
<code>show appliance usergroup</code>	List all User groups. Within each group, list each User assigned to the group, User type, and User ID.
<code>setup appliance usergroup <group_name> add-members</code> <code>setup appliance usergroup <group_name> delete-members</code>	In NMC, add or delete multiple Users in an existing group using the command line.
<code>destroy appliance usergroup</code>	Delete an existing User group. You cannot delete a default User group (root, staff, or admin).

About LDAP Users

The Lightweight Directory Access Protocol (LDAP) is a common protocol interface to Network Directory Services. LDAP provides widely deployed directory services such as Domain Name Service (DNS) and Network Information Service (NIS). DNS and NIS provide clients with information such as host IP addresses, User names, passwords and home directories. LDAP typically runs over TCP, and has the potential to consolidate existing Network Directory Services into a single global directory.

NexentaStor LDAP integration provides the following advantages:

- User's connection and authentication control
- Option to use User authentication in NFS-based environment
- Usage of ACL across heterogenous file services instead of POSIX permissions and attributes.
- Usage of SSL

See Also:

- [About ACLs](#)

Supported LDAP Servers

NexentaStor supports the following LDAP Servers:

- OpenLDAP
- OpenDJ
- Apache DS
- Red Hat 389 Directory Service
- Oracle Directory Services
- Microsoft Active Directory

Configuring the LDAP Server

You must configure an external LDAP server, before you configure it on NexentaStor . You must have the following information for this procedure:

- Groups, Users, and/or netgroups service descriptors
- Base Domain Name
- LDAP Server(s) IP address, or host name
- Proxy Domain Name
- Proxy password
- Security Certificate (generated on the LDAP server)

You can find all this information in the LDAP directory tree on your LDAP server.

❖ *To configure the LDAP server, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the **LDAP Client** panel, click **Add CA certificate**.
3. Type a name for the certificate.
4. Paste the certificate information in the Certificate field.
5. Click **Apply**.
6. In the **LDAP Client** panel, click **Configure**.
7. In the **Configure** window, select the following options:
 - **LDAP config type**
manual
 - **Groups Service Descriptor**
Type the group descriptors.

Example:

```
ou=groups,dc=nexenta,dc=local
```

- **Netgroup Service Descriptor**

Type the NFS networking group descriptors.

- **Credential Level**

Select the appropriate credential level according to the security settings on LDAP server.

- **Base DN**

Type the default search Base DN.

Example:

```
dc=nexenta,dc=local
```

- **LDAP Servers**

Type the IP address or host name of the LDAP server.

- **Authentication method**

Select the method of authentication.

- none
- simple
- sasl/CRAM-MD5
- sasl/DIGEST-MD5
- tls:simple

- **Proxy DN**

Type the client proxy DN, if needed.

Example:

```
cn=admin,dc=nexenta,dc=local
```

- **Proxy Password**

Type the client proxy password, if needed.

- **Users Service Descriptor**

Type the User service descriptors.

Example:

```
ou=users,dc=nexenta,dc=local
```

8. Click **Apply**.

If LDAP server is configured correctly, the LDAP Users are automatically uploaded to NexentaStor.

9. Click **Settings > Users** to verify the LDAP Users.

Example of the Account Entry on the LDAP Server

NexentaStor can work with various LDAP servers. In order to upload the Users to NexentaStor, the account entries on LDAP server must have the following format:

```
dn: uid=john,ou=users,dc=nexenta,dc=local
objectClass: inetOrgPerson
objectClass: posixAccount
objectClass: shadowAccount
uid: John
sn: Doe
givenName: John
cn: John Doe
displayName: John Doe
uidNumber: 10000
gidNumber: 5000
gecos: John Doe
loginShell: /bin/bash
homeDirectory: /home/john
```

Uploading an SSL Certificate

NexentaStor uses a self-signed SSL certificate for LDAP connections. You can change the SSL certificate if your company security policy requires to use a specific SSL certificate.

❖ *To upload an SSL Certificate, using NMV:*

1. Click **Settings > Network**.
2. In the **Network** pane, click **SSL/LDAP-Bind**.
3. Under **Certificate File**, click **Choose File**.
4. Select the file of the secure certificate.
5. Click **Open**.
6. Optionally, type the **Certificate Alias**.
7. Click **Save**.

❖ *To upload an SSL Certificate, using NMC:*

1. Type:

```
nmc:/$ setup network ssl-bind
System response:
```

```
Choose certificate source:
```

2. Select from the following options:
 - Remote SSL server
 - Local certificate file

- If you selected **Remote SSL server**:
 - 1) Type **SSL server host[:port]**.
 - 2) Type **Certificate Alias**.
- If you selected **Local certificate file**.
 - 1) Type the path to certificate file.
You must first copy the certificate file to the NexentaStor appliance.
 - 2) Type **Certificate Alias**.

Updating the SSL Certificate Information

You can update the information about the SSL Certificate, such as common name and number of days the certificate remains valid. These parameters will be reflected in the certificate information in your browser.

Before updating the information, complete tasks in [Uploading an SSL Certificate](#).

❖ *To update the SSL certificate information, using NMV:*

1. Click **Settings > Appliance**.
2. In the Administration pane, click **SSL Certificate**.
3. Under **Certificate File**, click **Choose File**.
4. Fill the following fields with the information from the SSL certificate:

Table 9-5: SSL Certificate Information Parameters

Country	Country name in 2 letter code. For example, US.
State/Province	Name of the State or Province. For example, California.
City/Locality	Name of the city or locality. For example, San Francisco.
Organization/Company	Name of the company or organization. For example, Nexenta Systems, Inc.
Department	Name of the organizational unit. For example, Storage, Technology.
Common Name	Hostname of the server. For example, nexenta.com. The Common Name must be equal to the fully qualified domain name by which the machine address is resolved over the network. However, it does not require to be equal to the local machine FQDN (that includes values stored in <code>/etc/nodename</code> and <code>/etc/resolv.conf</code>). In case you change this value, you may need to generate the new certificate.
Email Address	Email address.
Days	Number of days during which the certificate is valid.

5. Click **Save**.

❖ *To update the SSL certificate information, using NMC:*

1. Type:

```
nmc:/$ setup appliance init
```

System response:

```
Reconfigure? (y/n)
```

2. Type n.

System response:

```
Your primary interface is : e1000g0
Web GUI protocol          :
  HTTP  HTTPS
```

3. Select HTTPS.

4. Type new or leave default Web GUI port.

System response:

```
Starting self-signed SSL-certificate for NMV generation.
Your appliance has pre-generated SSL certificates. To improve security
it's recommended to generate new certificate. Generate new certificate?
(y/n)
```

5. Type n.

6. Follow the prompts to complete the required parameters.

For more information, see [SSL Certificate Information Parameters](#).

This page intentionally left blank

Managing the SCSI Target

This section includes the following topics:

- [About the NexentaStor SCSI Target](#)
- [Creating a Target Portal Group](#)
- [Creating an iSCSI Target](#)
- [Creating a Target Group](#)
- [About ZVOL](#)
- [Creating a ZVOL](#)
- [Viewing and Changing ZVOL Properties](#)
- [Destroying a ZVOL](#)
- [About LUN Mapping](#)
- [Creating a LUN Mapping](#)
- [About Secure Authentication](#)
- [Creating a Remote Initiator on a NexentaStor Appliance](#)
- [Configuring as NexentaStor iSCSI Initiator](#)
- [Creating an Initiator Group](#)

About the NexentaStor SCSI Target

A SCSI Target is a generic term used to represent different types of targets, such as iSCSI or Fibre Channel. The NexentaStor SCSI Target feature accesses all targets in the same manner. This allows the same zvol to be exported to any type of target, or to multiple targets simultaneously.

Configuring a target means making it available to the system. This process is specific to the type of target you configure.

You can use the SCSI Target software framework to create iSCSI targets that can be accessed over a storage network by iSCSI initiator hosts.

NexentaStor Enterprise Edition and the NexentaStor Trial Version contain a pre-installed SCSI Target 2.0 plug-in.

The NexentaStor SCSI Target 2.0 plug-in is an extension that provides the following features:

- Fully integrated into NexentaStor and provides the UI through NMV.
- Exposes one or more zvols as iSCSI LUNs over the SAN.

- Creates multiple iSCSI targets independent of the number of LUNs exposed through those targets.
- Creates mappings that allows specific sets of iSCSI initiators to see specific iSCSI LUNs through a specific set of iSCSI targets, instead of all of the iSCSI initiators having access to all of the iSCSI LUNs.
- Uses CHAP or Radius authentication for Initiators.
- Uses iSNS for device discovery.

Creating a Target Portal Group

Before creating a SCSI target, you assign a particular network interface (NIC) for iSCSI traffic. You do this by configuring a Target Portal Group (TPG).

Note: You cannot create a SCSI target before you configure at least one TPG.

❖ *To create a Target Portal Group, using NMV:*

1. Click **Data Management > SCSI Target**.
2. In the **iSCSI** panel, click **Target Portal Groups**.
3. In the **Manage iSCSI TPGs** window, click **here**.
4. In the **Create iSCSI TPG** window, type a name for the TPG, and also type at least one IP address of a NIC that you are going to use for iSCSI traffic, then click **Create**.

You can give any logical name to your TPG, such as **TPG_1**. The simplest example of an IP address is the IP address of the NexentaStor primary network interface.

❖ *To create a TPG, using NMC:*

1. Type:


```
nmc:/$ setup iscsi TPG
```
2. Follow the prompts to create the TPG.

Tip: You can enter the IP address without a port into the **Addresses** field. In that case, the portal uses the default value of **3260**. You can also enter multiple IPv4 portal addresses, if you plan to use more than one path for this TPG.

Creating an iSCSI Target

An iSCSI target is a software-only service that you can create as needed. You can provision one or more NexentaStor LUNs over iSCSI on a block level.

Tip: If you do not specify a name for the target, NexentaStor automatically assigns a target name.

- ❖ *To create an iSCSI target, using NMV:*
 1. Click **Data Management > SCSI Target**.
 2. In the iSCSI panel, click **Targets**.
 3. In the **Manage iSCSI Targets** window, click **here**.
 4. In the **Create iSCSI Target** panel, type a name or click **Create** to assign the target name automatically.
 5. Select the authentication type.
 6. Select other options, as needed.
- ❖ *To create an iSCSI target, using NMC:*
 1. Type:


```
nmc:/$ create iscsi target
```
 2. Follow the prompts and enter the target name, authentication, and other information.

Note:

You can also edit the other fields, such as Alias and a CHAP User/secret. The CHAP credentials that you specify during the iSCSI Target creation are only relevant for bidirectional CHAP. For non-bidirectional CHAP authentication, specify the CHAP parameters for the initiator on the Initiators page.

Creating a Target Group

You can create multiple iSCSI targets on the same NexentaStor appliance for a multi-path scenario. You can also unite these targets in logical groups to control the access of initiator hosts.

- ❖ *To create a target group, using NMV:*
 1. Click **Data Management > SCSI Target**.
 2. In the **SCSI Target** panel, click **Target groups**.
 3. In the **Manage Target Groups** window, click **here**.
 4. In the **Create New Target Group** panel:
 - Type a Group Name
 - Select targets.
 5. Click **Create**.

About ZVOL

A zvol is a virtual block device on a volume. In effect, it is a LUN that is remotely accessible through the protocols supported by SCSI Target plugins. The zvol can be managed, compressed, replicated, have snapshots taken of it, and so on.

Creating a ZVOL

To provision storage over iSCSI or FC, you must create a zvol on an existing volume.

❖ *To create a zvol, using NMV:*

1. Click **Data Management > SCSI Target**.
2. In the **ZVOLS** panel, click **Create**.
3. In the **Create a New ZVOL** window, fill in the required fields and click **Create**.

❖ *To create a zvol, using NMC:*

1. Type:

```
nmc:/$ setup zvol create
```
2. Follow the prompts to enter the volume name, size, and other properties

Table 10-1: ZVOL Properties

Name (NMV)	Name (NMC)	Description
Volume	zvol name	Pathname of a zvol.
Name		Consists of the name of a volume, and a unique zvol name, separate by forward slash (/).
Size	zvol size	Capacity of a zvol. For example: 2TB, 100MB, 500K. The capacity can be later expanded using the <code>grow</code> command in NMV or NMC.
Block Size	block size	Block size of a disk. The default value is 128K.
Compression	compression	Compression algorithm used for the zvol.
Initial reservation	Create a thin provisioned device?	Yes — Create a thin provisioned device that grows gradually to the maximum size of the zvol. No — Creates a zvol with the reserved entire size.
Description		Short description of the virtual block device (optional).
Deduplication		If enabled deletes redundant copies of data. Deduplication uses system RAM to store the deduplication hash table. Consult with the NexentaStor representative about the usage of deduplication before enabling.
Log Bias		Manages synchronous requests. The options are: <ul style="list-style-type: none"> • Latency — Sends data to a separate log device to proceed with low latency • Throughput — Writes data directly to the pool device at high performance and the write workloads spread.

Table 10-1: ZVOL Properties

Name (NMV)	Name (NMC)	Description
Number of copies		Number of copies stored for this dataset.
Sync		Controls synchronous requests: Standard — all synchronous requests are written to stable storage Always — every file system transaction is written and flushed to stable storage by system call return Disabled — synchronous requests are disabled

See Also:

- NexentaStor FC Plugin User Guide
- [Viewing and Changing ZVOL Properties](#)

Viewing and Changing ZVOL Properties

You can view and change the properties of a ZVOL.

The **Size** value must be a multiple of the **Block Size**.

Note: Use caution when reducing the size of a ZVOL. If the new size is smaller than the current amount of data in the ZVOL, some data might be lost.

❖ *To view and change the properties of a ZVOL, using NMV:*

1. Click **Data Management > SCSI Target**.
2. In the **ZVOLS** panel, click **View**.
3. In the **View ZVOLS** window, click the name of the ZVOL.
4. Modify the appropriate property field.
5. Click **Save**.

❖ *To view the properties of a ZVOL, using NMC:*

◆ **Type:**

```
nmc:/$ show zvol <vol_name/zvol_name> -v
```

❖ *To change the properties of a ZVOL, using NMC:*

1. Type:

```
nmc:/$ setup zvol <vol_name/zvol_name> property
```

2. Type the appropriate property name and press **Enter**.

Example:

```
nmc:/$ setup zvol <vol_name/zvol_name> property volsize
```

3. Follow the onscreen prompts to modify the value.

Destroying a ZVOL

You can destroy/delete a ZVOL.

❖ *To delete a ZVOL, using NMV:*

1. Click **Data Management > SCSI Target**.

2. In the **Zvols** panel, click **View**.

3. In the **View ZVOLs** window, select the appropriate ZVOL and click **Delete**.

❖ *To delete a ZVOL, using NMC:*

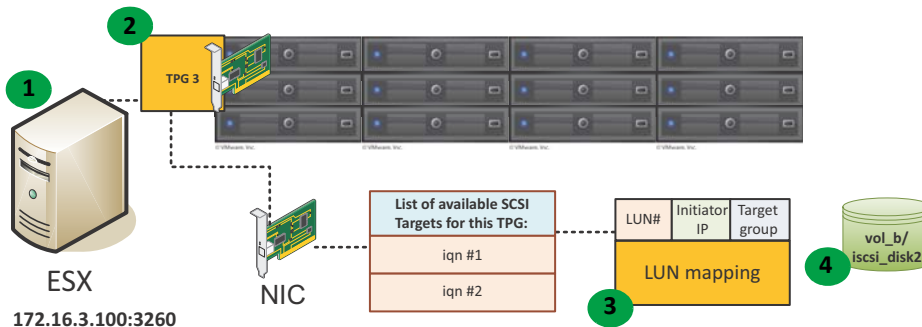
◆ Type:

```
nmc:/$ setup zvol <vol_name/zvol_name> destroy
```

About LUN Mapping

LUN mapping allows you to open particular LUNs to specific initiators and hide them from other initiators. You can use target and initiator groups to manage the access to specific zvols.

LUN mapping works as follows.



#	Description
1	Initiator connects to a specific IP address.
2	Initiator acquires the list of available SCSI Targets for the specific TPG.
3	Initiator acquires the list of LUNs available for the SCSI targets.
4	Initiator acquires the LUN mapping information and connects to the available LUN.

Creating a LUN Mapping

LUN mappings enable you to select which iSCSI targets export the current LUN and which initiators can see it. In the simplest configuration, all initiators can see all targets.

Note: The LUN Mapping feature only works in NMV.

❖ *To create a LUN mapping, using NMV:*

1. Click **Data Management > SCSI Target**.
2. In the **SCSI Target** panel, click **Mappings**.
3. In the **Manage Mappings** window, click **here**.
4. In the **Create New Mapping** dialog:
 - Select a zvol
 - Select an initiator group
 - Select a target group
 - Type a LUN number (optional)
 - Type a serial number (optional)

If you do not have an Initiator or Target group, you can create All-to-All mapping. This configures all initiators to see this zvol and configures all targets to expose it. You can add as many mappings as you need to configure the required access control list.

5. Click **Create**.

Tip: You can also share a zvol to all the hosts and initiators by marking the appropriate zvol and clicking **Share** on the **View Zvols** page. This creates a new zvol entry on the LUN Mappings page that all the initiators and all the targets can see.

About Secure Authentication

Setting up secure authentication is optional. You can use one of the following options to ensure that only the trusted hosts access the target.

- Unidirectional CHAP (see [Establishing Unidirectional CHAP Authentication](#))
- Bidirectional CHAP

Tip: Make sure that your hardware supports the chosen level of security.

Establishing Unidirectional CHAP Authentication

Challenge-Handshake Authentication Protocol (CHAP) is a scheme that the PPP protocol uses to authenticate the remote clients in the network.

Unidirectional CHAP is the most commonly used iSCSI security level. You can enhance data security and ensure that only authorized initiators access the data with unidirectional CHAP between a particular initiator and NexentaStor on a peer-to-peer model.

Unidirectional CHAP assumes that an initiator has its own secret, which you specify on the NexentaStor side. When an initiator connects to a target, SCSI Target checks the initiator's credentials and grants access to data. For every initiator logging into a NexentaStor iSCSI target with unidirectional CHAP enabled, you must create a Remote Initiator and set its CHAP secret.

❖ *To establish unidirectional CHAP authentication:*

1. On the iSCSI Initiator side, create a CHAP secret.

Example:

In Windows 7 click **Start > All Programs > Administrative Tools > iSCSI Initiator > Configuration**, click **CHAP** and specify the 12-character password.

2. On the NexentaStor side, perform the steps explained in [Creating a Remote Initiator on a NexentaStor Appliance](#) and specify the CHAP credentials.

You can leave the field **CHAP User** blank. If you do this, the initiator node name is assigned automatically

See Also:

- [Creating an iSCSI Target](#)

Establishing Bidirectional CHAP

Bidirectional CHAP provides a two-layer authentication protection. It requires that the target identifies an initiator, and also that the initiator identifies the target.

You can establish bidirectional CHAP to provide more secure authentication. Set up a CHAP User name and password on the target side by choosing the CHAP authentication method when you create the iSCSI target.

- ❖ *To establish bidirectional CHAP:*
 - ◆ Follow the steps in [Creating an iSCSI Target](#). In the **Authentication Method** drop-down menu, choose **CHAP**.

See Also:

- [Creating an iSCSI Target](#)

Creating a Remote Initiator on a NexentaStor Appliance

You can create a remote initiator to a NexentaStor appliance to create the initiator group and control the access of initiator hosts to a particular zvol. You can also use the remote initiator functionality to establish unidirectional CHAP authentication.

You can obtain the initiator's name in the settings of the host operating system.

Example:

In Windows 7 you can select the **Initiator Name** field in **Start >All Programs > Administrative Tools > Configuration**. The Initiator name format displays as `iqn.1991-05.com.microsoft:<pc-name>`

- ❖ *To add a remote initiator, using NMV:*
 1. Click **Data Management > SCSI Target**.
 2. In the **iSCSI** panel, click **Remote Initiators**.
 3. In the **Remote Initiators** window, click **You can create a new one here**.
 4. In the **Create iSCSI Remote Initiator** window, type the name of the initiator.
 5. Click **Create**.
- ❖ *To add a remote initiator, using NMC:*
 1. Type:


```
nmc:/$ setup iscsi initiator <parameter> <Initiator_Name>
```
 2. Press **Enter** to accept the default initiator name, or backspace to delete the default name and type a new name, then press **Enter**.
 3. (Optional) Type a CHAP name:


```
nmc:/$ setup iscsi initiator <parameter> <CHAPName>
```

Configuring as NexentaStor iSCSI Initiator

NexentaStor can act as an initiator to another NexentaStor appliance, or any other storage system, so that your local NexentaStor can see the LUNs from another storage system/appliance on the local NexentaStor appliance.

❖ *To set up NexentaStor as an Initiator to another Appliance, using NMC:*

1. Type:

```
nmc:/$ setup iscsi discovery discovery-address add
```

System response:

```
iSCSI IP address:
```

2. Type the IP address of the iSCSI target.

Note:

The iSCSI target can be another Nexenta iSCSI target portal, another NexentaStor appliance, or a third-party storage system.

3. Go to **Data Management > Data Sets** and create a new volume with the iSCSI disks that NexentaStor just discovered.

Creating an Initiator Group

You can use initiator groups to restrict the access of various initiators to NexentaStor targets and to certain data located in specified volumes.

❖ *To create an Initiator group, using NMV:*

1. Click **Data Management > SCSI Target**.

2. In the **SCSI Target** panel, click **Initiator Groups**.

3. In the **Manage Groups of Remote Initiators** window, click **here**.

4. In the **Create New Initiator Group** window:

5. in the field **Group Name**:

- Specify a custom group name
- Specify Additional Initiators (optional)

6. Click **Create**.

Managing Snapshots and Checkpoints

This section includes the following topics:

- [About Snapshots and Checkpoints](#)
- [Setting up a Periodic Snapshot Service](#)
- [Using Snapshots](#)
- [Using Checkpoints](#)
- [Saving Configurations](#)
- [Restoring Configurations](#)

About Snapshots and Checkpoints

A snapshot is a read-only point-in-time representation of a file system, volume or zvol. You can clone a snapshot to create an editable copy.

Snapshots allow safety across reboots and upgrades. You can create an almost unlimited number of snapshots because they do not require any additional storage. Snapshots are stored on the same disk as the source dataset. There are two types of snapshots:

- A snapshot of a folder, volume, or zvol is called a snapshot.
- A bootable snapshot of the appliance's operating system is called a checkpoint. A checkpoint is automatically created when you upgrade the base appliance software or install additional plugin modules.
 - A rollback checkpoint is a bootable snapshot that contains the system state before upgrade. If the upgrade fails, you can reboot the appliance and boot to the rollback checkpoint.
 - An upgrade checkpoint is a snapshot of the updated system which is activated by the administrator after the upgrade and reboot.

Checkpoints are visible in the GRUB boot menu. You can boot into any saved checkpoint.

See Also:

- [Editing NMS Properties](#)

Setting up a Periodic Snapshot Service

You can create an automatic periodic snapshot service of a folder, volume or zvol. You can create multiple AutoSnap services with a unique schedule for each dataset.

See Also:

- [Creating an Auto-Service](#)
- [Auto-Snap](#)

Using Snapshots

NexentaStor provides a number of options for working with snapshots.

Creating a Snapshot

You can create a snapshot for any folder before making any significant changes to the folder. Then, you can roll back the state of the folder, if needed.

❖ *To create a snapshot, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Snapshots** panel, click **Create**.
3. In the **Create New Snapshot** window, choose a dataset and type a name for the snapshot. You can also select the **Recursive** option.
4. Click **Create**.

In the **Summary Information: Snapshots** a new snapshot displays.

❖ *To create a snapshot, using NMC:*

◆ **Type:**

```
nmc:/$ create snapshot <vol_name>/<fol_name>@<snapshot_name>
```

Example:

```
nmc:/$ create snapshot vol1/fol1@snap1
```

System response:

PROPERTY	VALUE	SOURCE
name	vol1/fol1@snap1	
type	snapshot	-
creation	Tue Aug 28 15:06 2012	-
used	0	-
referenced	31K	-
compressratio	1.00x	-
devices	on	default
exec	on	default
setuid	on	default
xattr	on	default
version	5	-
utf8only	off	-
normalization	none	-
casesensitivity	mixed	-

nbmand	off	default
primarycache	all	default
secondarycache	all	default
defer_destroy	off	-
userrefs	0	-
mlslabel	none	default
nms:dedup-dirty	off	inherited from vol1/foll

Viewing Snapshot Summary Data

You can view summary information about snapshots.

❖ *To view the snapshots, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Snapshots** panel, click **Show**.

❖ *To view the snapshots, using NMC:*

1. Type:


```
nmc:/$ show snapshot
```
2. View the information about the snapshot.

Cloning a Snapshot

A clone of a snapshot is a readable, writeable copy of a snapshot. You can clone a snapshot and manage the clone separately on any storage appliance.

If you want to clone a snapshot that was created by an Auto-Sync, Auto-Snap, or Auto-Tier service, you must first remove the service label from the snapshot using the `take-ownership` command. This functionality is only available in NMC

❖ *To clone a snapshot, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Snapshots** panel, click **Show**.
3. In the **Summary Information: Snapshots** window, click the clone icon in the appropriate snapshot row.
4. In the **Clone snapshot** dialog:
 1. Select a folder for the clone
 2. Type a name for the clone
 3. Click **Clone**.

The clone is stored in the destination folder. View it by opening **Data Management > Data Sets** and clicking **Folders > Show**.

❖ *To clone a snapshot, using NMC:*

1. If the snapshot is created by an Auto-Sync, Auto-Snap, or Auto-Tier service, remove the service label by typing:

```
nmc:/$ setup snapshot <vol_name@snapshot_name> take-ownership
```

2. Type:

```
nmc:/$ setup snapshot <vol_name@snapshot_name> clone
```

3. Follow the prompts to complete the operation.

Recovering Files From a Snapshot

There are two methods for recovering files from a snapshot:

- Clone a snapshot in Windows or Illumos, and access the read and write clone of a snapshot (not available for zvol).
- Copy the contents of the hidden .zfs directory

❖ *To access a clone of a snapshot, using NMV:*

1. Share the folder using CIFS or NFS.
2. Click **Data Management > Data Sets**.
3. In the **Folders** panel, click **Show**.
4. In the **Edit Folder: FolderName** window, change the folder's **Snapshot visibility** property to **visible**.

❖ *To access a clone of a snapshot, using NMC:*

1. Type:

```
nmc:/$ setup folder <fol_name> property snapdir
```

2. In the snapdir field, type: `visible`

3. Log in to the client machine and map the shared folder.

If you use Windows, map the share as a network drive.

4. Access the shared folder on the client machine.

The default login and password for NexentaStor CIFS share is **smb/nexenta**.

5. Open `\.zfs\snapshot\<snapshot_name>`.

6. Copy the files in this directory to a new location.

See Also:

- [Cloning a Snapshot](#)
- [Sharing Folders](#)

Rolling Back to a Snapshot

You can roll back a folder, volume or zvol to a snapshot of its previous state.

❖ *To roll back a folder, volume or zvol to a specific snapshot, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Snapshots** panel, click **Show**.
3. In the **Summary Information: Snapshots** window, select the snapshot and click the **Rollback** icon.
4. In the confirmation dialog, click **OK**.

❖ *To roll back a folder, volume or zvol to a specific snapshot, using NMC:*

1. Type:

```
nmc:/$ setup folder <fol_name> snapshot <snapshot_name> rollback
```

2. Select **yes** to rollback recursively, or **no** to rollback non-recursively.
3. Type **y** to confirm the operation.

Example:

```
Rollback dataset 'voll/foll'? The operation will revert all changes since
the time snapshot 'voll/foll@snap1' was taken - that is, since Tue Aug
28 15:06 2012. Proceed?
(y/n) y
```

System response:

```
Rolled back dataset 'voll/foll' to 'voll/foll@snap1' dated Tue Aug 28
15:06 2012
```

Deleting a Snapshot

You can delete snapshots. Because of the copy-on-write mechanism of ZFS, deleting a snapshot might not result in additional free space because other snapshots might use the same blocks as the deleted snapshot.

Note: You cannot delete a snapshot that has clones. You must delete the clones first.

❖ *To delete a snapshot, using NMV:*

1. Click **Data Management > Data Sets**.
2. In the **Snapshots** panel, click **Show**.
3. If needed, use the **Filter** field to specify parameters of the snapshots to delete.
4. Select the desired snapshots and click **Delete**.

❖ *To delete a snapshot, using NMC:*

1. Type:

```
nmc:/$ destroy snapshot <snapshot_name>
```

2. Type **y** to confirm the operation.

Example:

```
nmc:/$ destroy snapshot 'vol1/foll@snap1'? (y/n) y
```

System response:

```
Snapshot 'vol1/foll@snap1' destroyed
```

Renaming a Snapshot

You can rename snapshots.

❖ *To rename a snapshot, using NMC:*

◆ Type:

```
nmc:/$ setup snapshot <vol_name@snapshot_name> rename
<new_snapshot_name>
```

Using Checkpoints

NexentaStor provides a number of services for managing checkpoints.

Creating a Checkpoint

You can create a system bootable checkpoint of the system configuration at a specific point in time.

❖ *To create a checkpoint, using NMC:*

1. Click **Settings > Appliance**.

2. In the **Upgrade Checkpoints** panel, click **Create**.

3. In the **Create a new checkpoint** window, type a name of the checkpoint and click **Save**.

In the **Checkpoints** window, a new checkpoint with type *rollback checkpoint* displays.

Tip: You can review various checkpoint parameters in the **Checkpoints** panel.

❖ *To create a checkpoint, using NMC:*

◆ Type:

```
nmc:/$ setup appliance checkpoint create
```

Rolling Back to a Checkpoint

You roll back to a NexentaStor checkpoint to restore the state of a system at a particular point in time.

❖ *To roll back your system to a checkpoint, using NMV:*

1. Click **Settings > Appliance**.
2. In the **Snapshots** panel, click **Show**.
3. In the **Upgrade Checkpoints** window, click **View**.
4. In the **Checkpoint** window, select the checkpoint and click **Activate**.
5. In the confirmation dialog, click **OK**.
6. Reboot NexentaStor.

❖ *To roll back your system to a checkpoint, using NMC:*

1. Type:

```
nmc:/$ setup appliance checkpoint restore
```
2. Follow the instructions onscreen to select the appropriate checkpoint and system volume to complete the rollback.

Viewing Existing Checkpoints

You can view the current list of checkpoints and information about each one.

❖ *To view the checkpoints, using NMV:*

1. Click **Settings > Appliance**.
2. In **Upgrade Checkpoints**, click **View**.
3. Double-click a checkpoint name to display more information about the checkpoint.

❖ *To view the checkpoints, using NMC:*

1. Display the list of checkpoints by typing:

```
nmc:/$ show appliance checkpoint
```
2. To display a specific checkpoint, type:

```
nmc:/$ show appliance checkpoint <checkpoint_name>
```

Changing the Number of Checkpoints Kept

By default, NexentaStor stores the eight most recent checkpoints. You can change this value. When the maximum number of checkpoints is reached, NexentaStor deletes the oldest checkpoint and saves the new one.

Note: You can manage the number of snapshots in the GRUB menu using the NMS properties `upgrade_menu_keep` and `upgrade_rootfs_keep`. This option is only available using NMC.

❖ *To change the number of checkpoints stored, using NMC:*

◆ Type:

```
nmc:/$ setup appliance nms property upgrade_rootfs_keep <number>
```

where `<number>` is the maximum number of checkpoints to store.

Permanently Saving a Checkpoint

You can mark a checkpoint so that it will not be deleted. You can also remove the mark so that the checkpoint will be deleted when NexentaStor reaches the maximum number of checkpoints to store.

This option is only available using NMC.

❖ *To permanently save a checkpoint, using NMC:*

◆ Type:

```
nmc:/$ setup appliance checkpoint <checkpoint_name> hold
```

❖ *To release the mark and allow the checkpoint to be deleted, using NMC:*

◆ Type:

```
nmc:/$ setup appliance checkpoint <checkpoint_name> release
```

Deleting a Checkpoint

You can delete a checkpoint, or multiple checkpoints, immediately to free space on the disk.

Note: Because of the copy-on-write mechanism of ZFS, deleting a snapshot might not result in additional free space because other snapshots might use the same blocks as the deleted snapshot.

❖ *To delete a checkpoint, using NMC:*

1. Display the list of checkpoints by typing:

```
nmc:/$ show appliance checkpoint
```

2. Type:

```
nmc:/$ destroy appliance checkpoint <checkpoint_name>
```


Saving Configurations

Use the `save` command to preserve the configuration of all NexentaStor settings. By default, NMS saves NexentaStor configurations as tarballs in `.user_config`. You can modify the location of saved configuration.

This functionality is only available in NMC.

- ❖ *To view the location of the saved configuration, using NMC:*

- ◆ Type:

```
nmc:/$ setup appliance configuration location
```

- ❖ *To save the NexentaStor configuration to the default directory, using NMC:*

- ◆ Type:

```
nmc:/$ setup appliance configuration save
```

- ❖ *To save the NexentaStor configuration to a specific folder, using NMC:*

- ◆ Type:

```
nmc:/$ setup appliance configuration save -F <directory-path>
```

The directory path can be relative or absolute.

- **Relative directory path**

An appliance-specific configuration (mailer, plug-ins, hostname settings, etc.) is saved on the spool. An auto-services configuration for each volume is saved on each volume.

- **Absolute directory path**

All configuration information is saved to the specified directory, and a sub-directory is created for each volume. Use this option to make a backup of the entire configuration.

Restoring Configurations

Use the `restore` command to reset the configuration of NexentaStor to the settings in the previously saved configuration file.

NexentaStor restores the latest saved configuration. You cannot select a configuration file to restore.

When restoring a configuration, you have multiple choices. Depending on your choice, NMS restores full configuration or the selected components.

Note: This functionality is only available in NMC.

- ❖ *To restore all or part of a saved auto-services configuration, using NMC:*
 - ◆ **Type:**

```
nmc:/$ setup appliance configuration restore
```

Selecting **all** restores all options listed, including information about Active Directory, ID mapping, network IP address, iSCSI targets, NFS and CIFS shares, share permission settings, and so on.
- ❖ *To restore the auto-services configuration of one volume only, using NMC:*
 - ◆ **Type:**

```
nmc:/$ setup appliance configuration restore -V <volume-name>
```
- ❖ *To restore an auto-services configuration from a specific folder, using NMC:*
 - ◆ **Type:**

```
nmc:/$ setup appliance configuration restore -F <directory-name>
```

The following table describes options for the configuration restore.

Table 11-1: Restore Configuration Options

all	Full appliance configuration
auto-scrub	Auto-scrub services configuration
auto-snap	Auto-snap services configuration
auto-sync	Auto-sync services configuration
auto-tier	Auto-tier services configuration
basic	Basic appliance settings, such as keyboard layout, locales, and so on.
hosts	Entries in the Internet host table.
interface	Configuration of network interfaces
iscsi	Settings of iSCSI Initiator and discoveries. This option does not restore CHAP passwords.
jbods	Settings of JBOD(s)

Table 11-1: Restore Configuration Options

<code>mailer</code>	Configuration of notification system including: <ul style="list-style-type: none"> • SMTP server • SMTP user name • SMTP Password • SMTP Timeout • SMTP Authentication • Email Addresses for general notifications NMS does not keep settings for the optional fields. Therefore, after restore NMS keeps the modified settings and does not roll them back to saved parameters.
<code>names</code>	The NexentaStor appliance host name and domain name.
<code>netshvc</code>	Restores statuses of network service, such as NFS, CIFS, RSYNC, and so on.
<code>nmc</code>	NMC configuration
<code>nms</code>	NMC configuration
<code>nmv</code>	NMC configuration
<code>routes</code>	Routing settings
<code>rsf-cluster</code>	HA cluster configuration
<code>runners</code>	Runners configuration
<code>scsitarget</code>	Configuration of SCSI targets
<code>shared-services</code>	Restores statuses of network services.
<code>shared-volumes</code>	Parameters of shared volumes
<code>shares</code>	Shared folders including CIFS, NFS, FTP, and so on. This command does not restore the status of the network service. Therefore, some of the shared folders may be unavailable. Run the <code>setup appliance configuration restore netshvc</code> command to restore the status.
<code>slotmap</code>	Slotmap configuration
<code>ssh-client</code>	SSH-client keys, the <code>known_host</code> file, and NMS bindings.
<code>ssh-server</code>	SSH server keys and the <code>authorized_keys</code> files.
<code>users</code>	User and group settings

Managing Auto-Services

This section includes the following topics:

- [About Auto-Services](#)
- [Replication Auto-Services](#)
- [Auto-Tier](#)
- [General Auto-Services](#)
- [Creating an Auto-Service](#)
- [Viewing an Auto-Service](#)
- [Editing an Auto-Service](#)
- [Starting and Stopping an Auto-Service](#)
- [Enabling and Disabling an Auto-Service](#)
- [Destroying an Auto-Service](#)

About Auto-Services

A NexentaStor auto-service is a service (function) that performs a specific task on a logical dataset. It creates an instance based on the schedule you define and the logical dataset chosen.

A logical dataset can contain any type, or combination of, collections of data, such as:

- zvol
- vdev
- volume
- folder
- virtual dataset
- physical dataset

NexentaStor provides two types of auto-services:

- Replication auto-services that perform synchronous and asynchronous replication and backup.
- General auto-services that perform other actions.

Replication Auto-Services

NexentaStor helps you to implement and maintain a robust backup strategy by simplifying the task of backing up data onto disks for primary and secondary storage, remote site replication, and archiving. This is useful for standard storage strategies as well as for tiered storage systems.

Comparing Auto-Tier and Auto-Sync

Auto-tier and Auto-Sync both replicate changed data from one NexentaStor appliance to another NexentaStor appliance.

The functions have several important differences.

Table 12-1: Comparison: Auto-Services

Feature Description	Auto-Tier	Auto-Sync
Main Benefit	Synchronizes source and destination folders according to replication schedule.	Synchronizes source and destination folders according to replication schedule.
Replication protocol	rsync	ZFS send/receive
Data replicated	New or modified files	New or modified blocks and metadata (snapshots, ACL, etc.)
Typical time needed per run	Hours	Less than Auto-Tier
Process for each execution	<ol style="list-style-type: none"> Determines the changed files. Sends the changed pieces of the files to the destination. 	<ol style="list-style-type: none"> Creates a snapshot on the source system. Uses ZFS to identify the newer blocks. Replicates only the changed blocks within the newer files to the target system.
System Types	Replicates to and from any type of file system	Replicates to and from ZFS file systems only
Number of Replications	Replicates to one file system	Replicates to multiple file systems

Auto-Tier

AutoTier is based on rsync. It is file-based so it replicates changed files within a specific source directory and sends them to a local or remote storage pool for common disk-to-disk backup. First, it replicates a snapshot. Then, it recursively searches through all files, and transfers only the newer files. It only transfers the changed blocks inside the newer files, including all nested sub-directories (except sub-directories specifically excluded).

By using snapshots on the destination, this tiered copy can have different retention and expiration policies and can be administered separately from the source file system.

You can create an Auto-Tier service for a folder, volume or zvol in order to synchronize the data between two computers or to create a back-up.

Auto-Tier provides the following advantages:

- Replication schedule
- Local-to local (L2L), local-to-remote (L2R), or remote-to-local (R2L) disk-to-disk replication of a volume, zvol or folder
- Multiple replication transports
- Source retention and expiration policy

Table 12-2: Auto-Tier Properties

Property	Definition
Name Custom Name	Optional parameter. User-friendly name for the logical dataset.
Direction	Direction of replication flow: <ul style="list-style-type: none"> • Local-to-Local • Local-to-Remote • Remote-to-Local <p>Note that in order to create L2R or R2L Auto-Tier services, two hosts must be previously SSH-bound to each other.</p>
Destination Folder	Depending on replication direction, the local or remote folder.
Replicate Content	Replicate content of the folder, but not the folder itself. This parameter means replicate only the source folder's sub-folders.
Periodic Interval	Consists of two parameters: Period and Frequency. Frequency values depend on period. Example: <ul style="list-style-type: none"> • Period: hourly; Frequency: 4 means run service every 4 hours. • Period: daily; Frequency: 05, means every 5th days.
Transport Protocol	There are several possible options: <ul style="list-style-type: none"> • RSYNC+ssh — Local host connects to remote using ssh and starts RSYNC session on the remote host. Then local and remote RSYNCs synchronize the directories. • RSYNC+TCP — Connection is established with remote RSYNC server. Replication performs through RSYNC-shared folders. • RSYNC+NFS — Remote nfs-share is mounted on the local appliance.
Remote source host	IP address or hostname of ssh-bound destination appliance.
Remote source folder	Destination folder. Example: <vol_name>/<fol_name>
Keep	Number of days to store snapshots on the local appliance.

Table 12-2: Auto-Tier Properties

Property	Definition
Trace level	The level of debug information displayed in the Auto-Tier log file.
Rate limit	Network traffic limitation, (Kb/s).
RSYNC FanOut/depth	Determines the depth of the Auto-Tier recursion.
Tiering snapshot	Starts the replication from any specified snapshot.
Exclude folders	Permits expelling folders from the replication stream. Example: You have <code>vol1/fol1</code> which contains sub-folders <code>a</code> , <code>b</code> , and <code>c</code> . To exclude <code>a</code> and <code>c</code> from the replication stream, type in the Exclude folders field (including the quotation marks): <code>"a c"</code> This excludes folders <code>a</code> and <code>c</code> from the replication stream.
RSYNC options	Extra RSYNC options. List of RSYNC arguments which you can use to start the RSYNC process. Type ' <code>man rsync</code> ' in NMC to get more information. For advanced Users only!
Service retry	Enables retrying the attempt to connect to the remote host in case of lost connectivity prior to expiration of the next scheduled interval.
Copy ACLs	In some cases Auto-Tier can copy a folder's ACL. This option works, if you use RSYNC+ssh protocol. In case of L2R or R2L replication, both hosts must be NexentaStor appliances.
Symlink as destination	Enables use of symlink as a destination folder. Caution: In this case, all files in the directory the symlink points to are overwritten.

Auto-Sync

AutoSync is a block-level replication service based on the send receive mechanism. Auto-Sync creates a snapshot on the source and it knows, directly from ZFS, and its creation times, what blocks to send to the other ZFS system. So, Autosync can operate between two ZFS systems.

Auto-Sync maintains a fully synchronized mirror of a given volume, folder, or zvol, on another NAS. The mirror includes all snapshots and file system properties. Auto-Sync uses the built-in ZFS snapshot capability to efficiently identify and replicate only changed blocks.

Auto-Sync creates a snapshot of a source folder and uses the Remote Replication protocol to send the snapshot to a destination folder. On subsequent iterations, Auto-Sync only sends the changes since the previous snapshot. Auto-Sync is built on the ZFS send/receive capability replication service. This minimizes traffic between the source and destination.

Auto-Sync is a part of NexentaStor base software.

Note: Auto-Sync may not be automatically installed with your version of NexentaStor. If it is not, go to **Settings > Appliance > Plugins** and click on the plugin to install it.

Auto-Sync provides the following advantages:

- Replication between NexentaStor and other ZFS-based operation systems
- Uses the native NexentaStor Remote Replication protocol
- Replicates one source to multiple destinations (multidestination)
- Replication of metadata
- Compresses the replication stream
- Sends only the changes between the two latest snapshots
- Resumes Auto-Sync in case of a replication failure

See Also:

- NexentaStor Auto-Sync User Guide

General Auto-Services

NexentaStor software provides various other auto-services.

Auto-Snap

Auto-Snap creates and destroys snapshots of a folder, volume or zvol on a schedule. You can use these snapshots to later restore the previous state of the folder and recover deleted files or sub-folders.

When you create a volume, a daily Auto-Snap service is automatically created for the volume.

Auto-Snap provides the following advantages:

- Automatically creates a point-in-time representation of a dataset
- Performs snapshots on a regular basis (daily, monthly, hourly)
- Creates snapshots of the root folder and its descendants
- Assists in creating a retention policy

Auto-Scrub

Auto-Scrub periodically performs volume cleaning operations. It is based on the ZFS capability of using end-to-end checksums to detect data corruption

Auto-Scrub checks each data block in a pool and verifies its consistency according to the checksum. If you use a redundant disk configuration, Auto-Scrub can restore the broken bits.

When you create a volume, an Auto-Scrub service is automatically created for the volume.

Note: Auto-Scrub consumes a lot of resources. It is most efficient to perform Auto-Scrub during a maintenance window.

Auto-Scrub provides the following advantages:

- Ensures data consistency
- Checks and repairs broken data blocks in mirrored configurations

Auto-Smart

Auto-Smart periodically checks the S.M.A.R.T. state of the disks.

See Also:

- *NexentaStor Auto-Smart Plugin User Guide*

Creating an Auto-Service

You can manually create any auto-service.

❖ *To create an auto-service, using NMV:*

1. Click **Data Management > Auto Services**.
2. In the corresponding auto-service panel, click **Create**.
3. In the **Create New Service** window, adjust the parameters of your auto-service and click **Create**.

❖ *To create an auto-service, using NMC:*

1. Type:

```
nmc:/$ setup <auto_service> <logical_dataset> create
```

Example:

```
nmc:/$ setup auto-sync vol5 create
```

2. Follow the prompts to specify the parameters for the auto service. These can include the interval, day, time, source, destination, dataset, or other parameters.

When NexentaStor combines the logical dataset with the schedule for the auto-service, it creates an Instance.

Viewing an Auto-Service

You can view the created auto-service to monitor its state and perform any required changes.

❖ *To view an auto-service, using NMV:*

1. Click **Data Management > Auto Services**.
2. In the appropriate auto-service panel, click **Show**.

3. To view the details of a specific Instance, double-click the Instance name for that service.

On this page you can start, stop, disable, enable, edit, delete and view the log of any auto-service.

- ❖ *To view auto-services, using NMC:*

- ◆ Type:

```
nmc:/$ show <auto_service> <instance>
```

Example:

```
nmc:/$ show auto-sync vol5-001
```

- ❖ *To view the schedule of all created services, using NMC:*

- ◆ Type:

```
nmc:/$ show <auto_service>
```

Editing an Auto-Service

You can edit an auto-service after its creation. Editable parameters might include service execution schedule, keep (store) days, recursive, exclude folders, trace level, and so on.

- ❖ *To edit an auto-service, using NMC:*

1. Click **Data Management > Auto Services**.
2. In the **Summary Information** panel, click the service to edit.
3. In the **Edit <service>** panel, perform the corresponding changes and click **Save**.

- ❖ *To edit an auto-service, using NMC:*

1. Type:

```
nmc:/$ setup <service> <instance> <property>
```

Example:

```
nmc:/$ setup auto-sync vol5-001 disable
```

2. Select the property and edit as required.

See Also:

- NexentaStor Auto-Sync User Guide

Starting and Stopping an Auto-Service

You can start or stop any auto-service using NMV on the **All Storage Services** summary page or on the specific auto-service summary page.

- When you start an auto-service, its status changes to **Online**. It immediately runs once and then runs again on its next scheduled execution.
- When you stop an auto-service, its status changes to **Offline**. It immediately stops, even if it is currently running, and it runs one more time on its next scheduled execution.

Note: If you need a service to not run on its next scheduled execution, you must disable the service.

❖ *To start or stop an auto-service, using NMV:*

1. Click **Data Management > Auto Services**.
2. In the auto-service panel, click **Show**.
3. In the **Summary Information** panel, select the service and click **Start** or **Stop**.

You can also click the name of a particular service to open the **Edit** window and select **Start** or **Stop**.

Note: You cannot start a service if the service is stopped once.

Enabling and Disabling an Auto-Service

You can temporarily turn off an auto-service, then re-enable it later. The service stops immediately when it is disabled, even if it is running.

After you enable a service, you must manually start the service before it runs.

❖ *To disable or enable an auto-service, using NMV:*

1. Click **Data Management > Auto Services**.
2. In the **Summary Information** panel, select the service and click **Disable** or **Enable**.

You can also click the name of a particular service to open the **Edit** window and select **Disable** or **Enable**.

Tip: You can disable or enable all auto-services by clicking **Disable All** or **Enable All** in the **All Storage Services** panel.

- ❖ *To disable or enable an auto-service, using NMC:*

- ◆ **Type:**

```
nmc:/$ setup <service> <instance> disable
```

or

```
nmc:/$ setup <service> <instance> enable
```

Example:

```
nmc:/$ setup auto-sync voll-005 disable
```

Destroying an Auto-Service

You can destroy an auto-service. This does not affect saved snapshots and backups created with these auto-services.

- ❖ *To delete an auto-service, using NMV:*

1. Click **Data Management > Auto Services**.
2. In the **Summary Information** panel, select the service and click **Delete Selected**. To delete an auto-service, using NMC:

- ◆ **Type:**

```
nmc:/$ setup <service> <instance> destroy
```

This page intentionally left blank

Monitoring NexentaStor Performance

This section includes the following topics:

- [About Performance Statistics](#)
- [Viewing Statistics](#)
- [Available Statistics](#)
- [Changing the Analytics Settings](#)
- [Viewing I/O Performance](#)
- [Using Performance Benchmarks](#)

About Performance Statistics

NexentaStor performance statistics are collected using DTrace. Statistics are displayed graphically using NMV and in tabular form using NMC.

NexentaStor analytics provide the following features:

- Displays representation of the DTrace real-time statistic in graphical view
- Analyzes performance on production systems in real time
- Generates performance profiles
- Analyzes performance bottlenecks
- Helps to troubleshoot problems by providing detailed views of the system internals

See Also:

- <http://dtrace.org/blogs/about>

Viewing Statistics

NMV displays statistics in a graphical format. NMC displays statistics in a character-based table format.

❖ *To view statistics, using NMV:*

1. Click **Analytics > Profiles**.
2. In the **Available Statistics** list, expand each desired section and select the metrics to include in the chart.

If you select more than one metric, make sure that the scale for each metric is compatible. If it is not, some charts might be difficult to read.


3. Click **Add Chart**.

The analytical system creates a chart that displays the selected metrics. The chart is updated in real time.

Note: NexentaStor combines all metrics that you select simultaneously into the same chart.

4. Arrange and close the charts as needed:

- To collapse a chart, click the up arrow in the chart's title bar.
- To delete a chart, click the X in the chart's title bar, or click **Remove** next to the chart name at the top of the **Profile Manager** panel.
- To switch a chart between bar chart display and line chart display, click the bar chart icon or line chart icon below the chart menu bar.



❖ *To view statistics, using NMC:*

1. Type:

```
nmc:/$ dtrace <analytic_type> <metric>
```

The metrics display and update in real time.

2. To stop the display and return to the command prompt, press **CTRL-C**.

To automatically stop the display and return to the command prompt after a certain time, include a number of seconds in the command.

Example:

To display statistics for 5 seconds and then automatically return to the command prompt, type:

```
nmc:/$ dtrace <analytic_type> <metric> 5
```

All NMC metrics include a ? option that displays brief help about the metric. Some NMC metrics include an `example` option that displays an example of the output with a longer explanation.

Available Statistics

You can generate charts according to various parameters.

Table 13-1: Available Statistics

Category	Statistic Name (NMV)
Virtual Memory Statistics	Page freed (Kbytes/Sec)
	Filesystem frees (Kbytes/Sec)
	Page reclaims (Kbytes/Sec)
	Page scan rate (pages/Sec)
	Filesystem page outs (Kbytes/Sec)
	Aggregated interrupts (number/Sec)
	Filesystem page ins (Kbytes/Sec)
	Aggregated system calls (number/Sec)
	Free memory (RAM) (Kbytes/Sec)
	Minor faults (Kbytes/Sec)
	Executable page outs (Kbytes/Sec)
	Anonymous page ins (Kbytes/Sec)
	Anonymous page outs (Kbytes/Sec)
	Swapped out lightweight processes (LWPs) (number/Sec)
	Executable page ins (Kbytes/Sec)
	Executable frees (Kbytes/Sec)
	Page ins (Kbytes/Sec)
	Anonymous frees (Kbytes/Sec)
	Page outs (Kbytes/Sec)
	Aggregated context switches (number/Sec)
Virtual memory (SWAP) free (Kbytes/Sec)	
Measure Aggregated CPU Utilization	Total CPU system mode (Percentage/Sec)
	Aggregated CPU idle mode (Percentage/Sec)
	Aggregated CPU busy mode (Percentage/Sec)
	Total CPU User mode (Percentage/Sec)
	Total CPU interrupt mode (Percentage/Sec)
Measure I/O throughput as generated by ZFS SPA sync function	

Table 13-1: Available Statistics

Category	Statistic Name (NMV)
	Average block size transfer through single SPA sync call (Kbytes/Sec)
	Calculated I/O ZFS SPA throughput (Mbytes/Sec)
	Total Megabytes transferred through single SPA sync (Mbytes/Sec)
	Number of milliseconds spent in SPA sync function (Milliseconds/Sec)
Storage I/O access pattern	
	Minimum I/O event size (Bytes/Sec)
	Percentage of events of a sequential nature (Percentage/Sec)
	Average I/O event size (Bytes/Sec)
	Total kilobytes read during interval (KBytes/Sec)
	Maximum I/O event size (Bytes/Sec)
	Percentage of events of a random nature (Percentage/Sec)
	Total kilobytes written during interval (KBytes/Sec)
	Number of IO events (Count/Sec)
TCP/IP Mib statistics	
	TCP bytes received out of order (Bytes/Sec)
	TCP bytes sent (Bytes/Sec)
	TCP bytes received (Bytes/Sec)
	TCP bytes received duplicated (Bytes/Sec)
	TCP bytes retransmitted (Bytes/Sec)
Number of interrupts aggregated by CPU	
	Number of Interrupts (number/Sec) CPU #

The following statistics are available using NMC.

COMSTAR

- fcwho
- iscsiio
- iscsirwlat
- iscsisnoop
- iscsiwho

CPU

- cputypes
- cpuwalk
- dispqlen
- hotkernel
- intbycpu
- intoncpu
- inttimes
- loads
- runocc
- xcallsbypid

IO

- bitesize
- diskhits
- fspaging
- fsrw
- hotspot
- iofile
- iofileb
- iopending
- iosnoop
- iotop
- rfileio
- rfsio
- rwsnoop
- rwtop
- seeksize
- statsnoop
- vopstat

Locks

- lockbydist
- lockbyproc

Memory

- anonpgpid
- minfbypif
- minfbyproc
- pgpginbypid
- pgpginbyproc
- swapinfo
- vmbypid
- vmstat
- vmstat -p
- xvmstat

Miscellaneous


- dtruss
- dvmstat
- execsnoop
- opensnoop
- statsnoop

Network

- icmpstat
- tcpsnoop
- tcpstat
- tcpwdist
- udpstat

Using Chart Profiles

You can save and load the charts later, with the specific parameters, to ease the monitoring of the appliance performance.

- ❖ *To save, load and delete custom charts:*
 - ◆ Use the   and  icons in the **Profile Manager**.

Removing the Custom Chart

You can delete unused charts from the **Chart Title** list.

- ❖ *To delete a chart:*
 - ◆ Click **Remove**.

Changing the Analytics Settings

You can change the period for the analytic to refresh.

- ❖ *To change the analytic refresh period, using NNV:*
 1. Click **Analytics > Settings**.
 2. In the **Analytics Settings** panel, in the **Analytics_refresh_interval** field, type the new value and click **Save**.

Viewing I/O Performance

You can view the real-time I/O performance of a data volume. This includes the allocated and free capacity, read/write operations, and bandwidth read/write.

- ❖ *To view the I/O performance of a volume, using NMC:*

- ◆ **Type:**

```
nmc:/$ show volume <vol_name> iostat
```

Using Performance Benchmarks

NexentaStor provides extensions that allow you to run stressful I/O and networking operations, and display the results. The extensions are provided as plugins.

You can install the following benchmarks:

- `bonnie-benchmark`
A filesystem measurement tool that runs a number of tests to determine the performance of a filesystem during the read and write operations.
- `iozone-benchmark`
A filesystem benchmark tool that measures the performance of read and write filesystem operations.
- `iperf-benchmark`
A network performance measuring tool that can emulate TCP and UDP data streams and measure their bandwidth performance.

See Also:

- [Running Bonnie Benchmarks](#)
- [Running IOzone Benchmarks](#)
- [Running Iperf Benchmarks](#)

Running Bonnie Benchmarks

The bonnie benchmark measures the performance of file systems by running a number of tests. You can run the test for a volume or a folder.

The bonnie-benchmark plugin runs the following tests:

- Sequential Write
- Sequential Read
- Random seek

You must install the `bonnie-benchmark` plugin before you can run the tests.

To run a test, the dataset should have enough free space.

This functionality is only available in NMC.

❖ *To run a bonnie benchmark, using NMC:*

1. Type:

```
nmc:/$ setup volume <volname> benchmark run bonnie-benchmark
```

System response:

```
WRITE      CPU    REWRITE   CPU    READ      CPU    RND-SEEKS
 162MB/s   8%     150MB/s   6%     188MB/s   9%     430/sec
 158MB/s   7%     148MB/s   8%     184MB/s   7%     440/sec
-----
 320MB/s  15%    298MB/s  13%    372MB/s  16%    870/sec
```

The following table describes options that you can specify for bonnie benchmark:

Table 13-2: Bonnie Benchmark Options

Option	Description	Default value
-p <numpros>	Number of process to run.	2
-b <blocksize>	Block size to use	32KB
-s	No write buffering. Run <code>fsync</code> after every write.	
-q	Quick mode.	

Example:

```
nmc:/$ setup volume <volname> bonnie-benchmark -p 2 -b 8192
```

See Also:

- [Running IOzone Benchmarks](#)
- [Running Iperf Benchmarks](#)
- <http://www.textuality.com/bonnie>

Running IOzone Benchmarks

The `iozone-benchmark` measures the performance of file systems by running filesystem operations including read, write, re-read, re-write, read backwards read strided, random read, and so on.

You must install the `bonnie-benchmark` plugin before you can run the tests.

To run the IOzone benchmark test, the dataset should have enough free space.

This functionality is only available in NMC.

❖ *To run an iozone benchmark, using NMC:*

1. Type:

```
nmc:/$ setup benchmark run iozone-benchmark
```

The following table describes options that you can specify for `iozone-benchmark`:

Table 13-3: Bonnie Benchmark Options

Option	Description	Default value
-p <numpros>	Number of process to run.	2
-b <blocksize>	Block size to use	32KB
-s	Max file size in MB to test.	auto
-q	Quick mode.	

See Also:

- [Running Bonnie Benchmarks](#)
- [Using Performance Benchmarks](#)
- <http://www.iozone.org>

Running Iperf Benchmarks

The `iperf-benchmark` plugin measures network performance by emulating TCP and UDP data streams. To run a test, you must have two NexentaStor appliances bound together using SSH and the `iperf-benchmark` plugin installed on both appliances.

This functionality is only available in NMC.

❖ *To run an Iperf benchmark, using NMC:*

1. Run the benchmark in server mode on one appliance by typing:

```
nmc:/$ setup benchmark run iperf-benchmark -s
```

System response:

```
-----
Server listening on TCP port 5001
TCP window size:  128 KByte (default)
-----
```

2. Run the benchmark in client mode on other appliance by typing:

```
nmc:/$ setup benchmark run iperf-benchmark -c
```

System response:

```
Remote appliance : <remote_ip>
-----
Client connecting to <remote_ip>, TCP port 5001
TCP window size:  256 KByte
-----
[ 5] local <local_ip> port 42851 connected with <remote_ip> port 5001
[ 4] local <local_ip> port 49550 connected with <remote_ip> port 5001
[ 3] local <local_ip> port 62744 connected with <remote_ip> port 5001
[ ID] Interval      Transfer    Bandwidth
[ 5]  0.0- 3.0 sec   203 MBytes  567 Mbits/sec
```

```

[ ID] Interval      Transfer      Bandwidth
[  3] 0.0- 3.0 sec    187 MBytes    523 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  4] 0.0- 3.0 sec    282 MBytes    789 Mbits/sec
[SUM] 0.0- 3.0 sec    672 MBytes    1.88 Gbits/sec
[ ID] Interval      Transfer      Bandwidth
[  3] 3.0- 6.0 sec    196 MBytes    547 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  4] 3.0- 6.0 sec    229 MBytes    641 Mbits/sec
[ ID] Interval      Transfer      Bandwidth
[  5] 3.0- 6.0 sec    278 MBytes    776 Mbits/sec
[SUM] 3.0- 6.0 sec    703 MBytes    1.96 Gbits/sec

```

You can run the test using basic settings or specify additional parameters.

The following table describes options for Iperf benchmarks.

Table 13-4: Iperf Benchmark Options

Options	Description	Default value
-s	Run <code>iperf-benchmark</code> in server mode on this appliance and in client mode on the remote appliance.	N/A
-c	Run <code>iperf-benchmark</code> in client mode on this appliance in server mode on the remote appliance.	N/A
-P	Number of parallel client threads to run	3
-i	Number of seconds between periodic bandwidth reports.	3 sec
-l	Length of buffer to read or write	128K
-w	The size of TCP window (socket buffer size).	256KB
-t	Total time in seconds to run the benchmark.	30 sec
hostname	If you run <code>iperf-benchmark</code> as client, you can specify the hostname or IP address of the Iperf server (optional).	N/A

Example:

```
nmc:/$ run benchmark iperf-benchmark <hostname> -P 5 -i 10 -l 8k -w 64k
```

See Also:

- [Running Bonnie Benchmarks](#)
- [Running IOzone Benchmarks](#)
- <http://iperf.sourceforge.net>

This page intentionally left blank

Monitoring SNMP Traps

This section includes the following topics:

- [About NexentaStor SNMP Implementation](#)
- [Using NexentaStor as SNMP Agent](#)
- [Using NexentaStor as SNMP Manager](#)

About NexentaStor SNMP Implementation

NexentaStor supports Simple Network Management Protocol (**SNMP**), a standard Internet protocol for managing devices over an IP network.

You can use NexentaStor to receive SNMP traps from other devices, or with any third-party SNMP management software, to monitor SNMP traps on a NexentaStor appliance. SNMP traps help to detect conditions that may require administrative action.

See Also:

- http://en.wikipedia.org/wiki/Simple_Network_Management_Protocol

Using NexentaStor as SNMP Agent

You can use any third-party SNMP monitoring tool to monitor NexentaStor SNMP traps. This section describes how to configure an SNMP Agent on NexentaStor.

To monitor traps on the NexentaStor appliance, see [Using NexentaStor as SNMP Manager](#).

About SNMP Traps

SNMP traps are warnings, of varying degrees of severity, that NexentaStor `snmpd` generates for specified events. SNMP traps are defined by MIB files. An SNMP trap is generated when the condition specified in a MIB file is met on the NexentaStor appliance.

About NexentaStor MIBs

A Management Information Base (MIB) is a text file that describes the structure of management data of the device subsystem. A MIB file contains the hierarchy of Object Identifiers (OIDs). An OID contains a variable that SNMP can read and translate to human readable form.

NexentaStor 4.x includes the following MIB files:

- NEXENTA-MIB.txt
- NEXENTA-ZFS-MIB.txt

All MIB files are stored in `/etc/net-snmp/snmp/mibs` on the NexentaStor appliance.

Note: Your SNMP monitoring tool may require the following MIB files:

- NETWORK-SERVICES-MIB.txt
- SNMP-FRAMEWORK-MIB.txt

Copy and load these files from `/etc/net-snmp/snmp/mibs` into your SNMP monitoring tool.

Using the NexentaStor MIBs

To monitor the NexentaStor performance using a third-party SNMP manager, import the NexentaStor MIB files to the corresponding SNMP monitoring tool.

Configuring the SNMP Agent

To monitor the NexentaStor failure events using a third-party SNMP Manager, configure the SNMP Agent on the NexentaStor appliance.

To complete this task, you must know the IP address of the SNMP management software that is installed in your environment.

❖ *To configure the SNMP agent, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the **SNMP Agent** panel, click **Configure**.
3. In the **SNMPv3 agent user** field, type the name of a user.
You can specify any user name. You do not need to specify a NexentaStor user.
4. In the **SNMPv3 agent password**, type the password for the SNMP user.
The SNMP agent user name and password are stored in the `snmp` configuration file on the NexentaStor appliance.
5. In the **Destination of SNMP-traps** field, type the IP address of SNMP management software.
6. Click **Apply**.

Note: Use your monitoring tool for further SNMP configuration.

❖ *To configure the SNMP agent, using NMC:*

1. Type:

```
nmc:/$ setup network service snmp-agent configure
```

2. Configure the parameters listed in the [SNMP Agent Parameters](#) table.

Table 14-1: SNMP Agent Parameters

Parameter	Description
Read-only community string	<p>The <code>read-only community string</code> parameter acts as a password, and enables the access to the SNMP device. The default value is <code>public</code>. Most of the network equipment is shipped with a default setting of <code>public</code>.</p> <p>You can leave the default setting, or modify it according to security requirements of your network environment.</p> <p>For simplicity, leave the default value.</p> <p>For more information, see the documentation for your SNMP monitoring tool.</p>
Read-write community string	<p>The <code>read-write community string</code> parameter enables you to modify writable MIB objects. The default value is undefined. However, you may want to change this value to <code>private</code> or leave it unmodified according to your security requirements.</p> <p>Do not set this value to <code>public</code>.</p> <p>For simplicity, leave the default value.</p> <p>For more information, see the documentation for your SNMP monitoring tool.</p>
User name for SNMPv3 agent	<p>Type the name of a user.</p> <p>You can specify any user name. You do not need to specify a NexentaStor user.</p>
Password for SNMP agent	<p>Type the password for the SNMP user.</p> <p>The SNMP agent user name and password are stored in the <code>snmp</code> configuration file on the NexentaStor appliance.</p> <p>The password must contain at least 8 characters.</p>
Destination of SNMP traps	Type the IP address of the SNMP monitoring tool.
System location	Optional commentary describing the location of your SNMP monitoring tool.
System contact	Optional descriptive commentary that specifies contact information.

Using NexentaStor as SNMP Manager

You can manage SNMP traps from the NexentaStor appliance using the SNMP Manager.

Enabling SNMP Manager

By default, SNMP Manager is disabled. Enable the SNMP Manager before starting to use it.

- ❖ *To enable SNMP Manager, using NMV:*
 1. Click **Settings > Misc. Services**.
 2. In the **SNMP Manager** panel, click **Enable**.

3. Click **OK** in the confirmation dialog.
 4. Proceed to [Configuring SNMP Trap Daemon](#).
- ❖ *To enable SNMP Manager, using NMC:*
 - ◆ Type:


```
nmc:/$ setup network service snmp-manager enable
```

Configuring SNMP Trap Daemon

Before proceeding, verify that you completed the tasks described in [Enabling SNMP Manager](#).

Configure general SNMP trap daemon parameters such as listening ports and logs and traps handling.

- ❖ *To configure SNMP Trap Daemon, using NMV:*
 1. Click **Settings > Misc. Services**.
 2. In the **SNMP Manager** panel, click **Configure**.
 3. Configure the fields described in the [SNMP Trap Daemon Settings](#) table.
 4. Click **Apply**.
- ❖ *To configure SNMP Trap Daemon, using NMC:*
 1. Type:


```
nmc:/$ setup network service snmp-manager configure
```
 2. Configure the settings described in the [SNMP Trap Daemon Settings](#) table.

Table 14-2: SNMP Trap Daemon Settings

Parameter	Description
Listening Address	<p>Defines the list of listening IP addresses on the NexentaStor appliance that will receive incoming SNMP notifications.</p> <p>Type a value in the following format: [<transport>:]<address></p> <ul style="list-style-type: none"> • transport — a transport protocol. Specify one of the following: <code>udp</code>, <code>tcp</code>, <code>unix</code>, <code>ipx</code>, <code>pvc</code>, <code>udp6</code>, <code>tcp6</code>, <code>ssh</code>, <code>dtlsudp</code>. • address — the <code>address</code> parameter format depends on the selected <code>transport</code>.
Do not retain notification logs	Disables collecting logs related to the SNMP Trap Daemon activity. The default value is yes.
Do not log traps	Disables logging of all notifications related to SNMP Trap Daemon.

Parameter	Description
Disable authorization	Disables access control verification, and enables all notifications. If you set this property to no, you must configure the list of authorized communities or users. Otherwise, no data will be transmitted. See Configuring Authentication .

See Also:

- [Configuring Authentication](#)
- [Creating SNMPv3 Users](#)

Creating SNMPv3 Users

Before proceeding, verify that you completed the tasks described in [Configuring SNMP Trap Daemon](#).

The SNMPv3 protocol uses username/password authentication. If you configured the SNMP authentication security level as `auth-Priv` or `authNoPriv` you must create at least one username/password pair.

❖ *To configure an SNMPv3 user, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the **SNMP Manager** panel, click **Configure SNMPv3 Users**.
3. In the **Operation to execute** field, select **create**.
4. Fill the corresponding fields, as described in [SNMPv3 User Configuration](#) table.
5. Click **Apply**.

❖ *To configure an SNMPv3 user, using NMC:*

1. Type:


```
nmc:/$ setup network service snmp-manager manage_users create
```
2. Type in the parameters as described in [SNMPv3 User Configuration](#) table.
3. OK.

Table 14-3: SNMPv3 User Configuration

Name in NMV	Name in NMC	Description
SNMPv3 trap/inform listener user	User name	Name of the SNMP user. The name should not contain any special characters, excluding '-' and '_'
User password	Password for SNMPv3 user	The user password must contain at least 8 characters.

Name in NMV	Name in NMC	Description
SNMPv3 encryption passphrase	Encryption passphrase	If you use the <code>authPriv</code> security level, specify an encryption passphrase. The passphrase must contain at least 8 characters. If you do not use encryption, leave this field empty.
SNMP engineID of remote agent	SNMP engineID of remote agent	An identifier in hex number format (0x1234..) that specifies the remote agent's engineID. If you leave the field empty, the engineID will be generated automatically.

See Also:

- [Removing SNMPv3 Users](#)

Configuring Authentication

Before proceeding, verify that you completed the tasks described in [Configuring SNMP Trap Daemon](#).

If you use SNMPv3, verify that you have created SNMPv3 users as described in [Creating SNMPv3 Users](#).

You can specify the types of messages that NexentaStor processes. If you configure the list of messages, NexentaStor will ignore other messages.

❖ *To configure SNMP Trap Daemon, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the **SNMP Manager** panel, click **Configure Authentication**.
3. Configure the fields described in the [SNMP Authentication Parameters](#) table.
4. Click **Apply**.

❖ *To configure SNMP Trap Daemon, using NMC:*

1. Type:


```
nmc:/$ setup network service snmp-manager configure_auth
```

2. Configure the settings described in the [SNMP Authentication Parameters](#) table.

Table 14-4: SNMP Authentication Parameters

Parameter	Description
Auth community string	<p>Specifies a string that defines the list of communities and traps that are authorized to be performed on the appliance. In SNMP, the concept of community string is similar to a user ID. Typically, SNMPv1 and SNMPv2 equipment has a predefined community string <code>public</code>. You can modify this setting as needed.</p> <p>This field is only used for SNMPv2 devices.</p> <p>Type a string in the following format:</p> <pre><type> <community> [<source> [<OID> -v <view>]]</pre> <ul style="list-style-type: none"> • <code>type</code> — specify the type of action: <ul style="list-style-type: none"> • <code>log</code> — add the detail of notification to a log file. • <code>net</code> — forward notification to another host. • <code>execute</code> — send the details about the trap to a handler program. • <code>community</code> — specify the name of community. For example: <code>public</code>, <code>private</code>. • <code>source</code> — specify that the configuration should only apply to notifications received from particular sources. This parameter is optional.
Auth user string	<p>Specifies a string that defines the list of users and traps that are authorized on the appliance. By default, authenticated requests are accepted.</p> <p>This field is only used for SNMPv3 devices.</p> <p>Type a string in the following format:</p> <pre><type> <user> [<level> [<OID> -v <view>]]</pre> <ul style="list-style-type: none"> • <code>type</code> — specify the type of action: <ul style="list-style-type: none"> • <code>log</code> — add the detail of notification to a log file. • <code>net</code> — forward notification to another host. • <code>execute</code> — send the details about the trap to a handler program. • <code>level</code> — specify the SNMP security level: <ul style="list-style-type: none"> • <code>authPriv</code> — use authentication with encryption • <code>authNoPriv</code> — use authentication without encryption. • <code>noAuth</code> — do not use authentication • <code>priv</code> — use encryption. • <code>OID</code> (or <code>-v view</code>)— specify particular notifications.

See Also:

- [Configuring SNMP Trap Handlers](#)

Configuring SNMP Trap Handlers

Before proceeding, verify that you completed the tasks described in [Configuring SNMP Trap Daemon](#).

When NexentaStor receives an SNMP trap, it either logs or forwards the trap to another host. You can specify a host to which you want to forward SNMP traps. You can also create a trap handler that will execute a specific action upon receiving an SNMP trap.

Before configuring an SNMP Trap Handler, place the program that you want the script to execute to in the following directory: `/var/lib/nza/snmp-scripts`

❖ *To configure SNMP trap handlers, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the **SNMP Manager** panel, click **Configure Trap Handling**.
3. In the **Trap Handlers** field, type a script.
See the [SNMP Trap Handler](#) table.
4. Click **Apply**.

❖ *To configure SNMP trap handlers, using NMC:*

1. Type:

```
nmc:/$ setup network service snmp-manager manage_trap_handling
```
2. Select **ADD**.
3. Type the script.
4. Select **OK**.
5. Type the script as described in the [SNMP Trap Handler](#) table.

Table 14-5: SNMP Trap Handler

Parameter	Description
Trap Handler	<p>Defines a script that invokes an action in response to receiving a notification with a specific OID.</p> <p>Type a script in the following format:</p> <pre><OID> default <program> [ARGS ...]</pre> <ul style="list-style-type: none"> • OID — the name of an OID. If you type <code>default</code>, the action will be performed upon receiving any SNMP trap. • program — the trap handler script name located at <code>/var/lib/nza/snmp-scripts</code>. • args — optional arguments for the script.

See Also:

- [Configuring SNMP Forwarding](#)

Configuring SNMP Forwarding

Before proceeding, verify that you completed the tasks described in [Configuring SNMP Trap Daemon](#).

You can forward SNMP traps to a specific destination.

❖ *To configure SNMP trap handlers, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the **SNMP Manager** panel, click **Configure Trap Handling**.
3. In the **Trap Forwards** field, type a script.
See the [SNMP Forwarding](#) table.
4. Click **Apply**.

❖ *To configure SNMP trap handlers, using NMC:*

1. Type:

```
nmc:/$ setup network service snmp-manager manage_trap_handling
```
2. Select OK.
System response:
Trap forwards:
3. Select ADD.
4. In the Type a script as described in the [SNMP Forwarding](#) table.
5. Select OK.

See Also:

- [Configuring SNMP Trap Handlers](#)

Table 14-6: SNMP Forwarding

Property	Description
Trap forwards	<p>Defines a script that forwards specific SNMP traps.</p> <p>Type a script in the following format:</p> <pre><OID> default <destination></pre> <ul style="list-style-type: none"> • <code>OID</code> — the name of an OID. If you type <code>default</code>, the action will be performed upon receiving any SNMP trap. • <code>destination</code> — destination host for trap forwarding. [<transport>:]<address> • <code>transport</code> — a transport protocol. Specify one of the following: <code>udp</code>, <code>tcp</code>, <code>unix</code>, <code>ipx</code>, <code>pvc</code>, <code>udp6</code>, <code>tcp6</code>, <code>ssh</code>, <code>dtlsudp</code>. • <code>address</code> — the address parameter format depends on the selected transport.

Removing SNMPv3 Users

You can delete an existing SNMPv3 user if you do not need the user anymore.

❖ *To delete an SNMPv3 user:*

1. Click **Settings > Misc. Services**.
2. In the **SNMP Manager** panel, click **Configure SNMPv3 Users**.
3. In the **Operation to execute** field, select **remove**.
4. Select a user to delete.
5. Click **Apply**.

❖ *To configure an SNMPv3 user, using NMC:*

◆ **Type:**

```
nmc:/$ setup network service snmp-manager manage_users remove
```

See Also:

- [Additional SNMP Manager Operations](#)
- [Configuring SNMP Trap Daemon](#)

Modifying snmptrapd.conf

You can modify the snmptrapd.conf file as needed. To modify the file, use standard SNMP methods. The snmptrapd.conf file remains empty before you configure the SNMP Trap Daemon.

See [Configuring SNMP Trap Daemon](#)

Warning: This topic is for advanced SNMP users. Do not modify the configuration file if you are not familiar with SNMP.

This functionality is only available in NMC.

❖ *To modify snmptrapd.conf, using NMC:*

1. **Type:**

```
nmc:/$ setup network service snmp-manager edit-settings
```

The snmptrapd.conf opens in vim editor.

2. Modify the file as needed.
3. Save changes by typing:

```
:wq!
```

See Also:

- [Additional SNMP Manager Operations](#)
- [Configuring SNMP Trap Daemon](#)

Additional SNMP Manager Operations

The following table lists additional commands that you can execute using SNMP Manager in NMC.

❖ *To run an SNMP manager command:*

1. Type:

```
nmc:/$ setup network service snmp-manager
```

2. Select a sub-command from the list.

Table 14-7: Additional SNMP Manager Operations

Command	Description
clear	Clear SNMP Manager faults.
confcheck	Verify the SNMP Manager service status
disable	Disable the SNMP Trap Daemon.
log	View the SNMP Trap Daemon log file.
restart	Restart the SNMP Trap Daemon service.
show	View information about the SNMP Manager.
unconfigure	Re-set all SNMP Manager parameters to default values and disable the service.

See Also:

- [Modifying snmptrapd.conf](#)
- [Configuring SNMP Trap Daemon](#)

Managing Nexenta Management Server

This section includes the following topics:

- [About Nexenta Management Server \(NMS\)](#)
- [Multi-NMS](#)
- [Editing NMS Properties](#)

About Nexenta Management Server (NMS)

The Nexenta Management Server is a service that manages all of the operations that a User needs to perform. The NMS is a layer between the Illumos kernel and its clients.

NMS runs in the background as a daemon. You cannot directly manage NMS. However, you can fine tune its performance.

NMS automatically performs the following tasks:

- Processes all requests from clients (NMV, NMC, and other client applications)
- Logs events
- Caches information about NexentaStor objects for faster access.
- Executes the implied logic
- Extends NexentaStor capabilities with plugins
- Controls NexentaStor remotely

Multi-NMS

Multi-NMS allows you to create several NMS instances. This facilitates work when you have multiple NMS clients. You can use single or multi-NMS mode depending on how much processing is involved with a single request or how many requests a server has to handle.

Multi-NMS is enabled by default and starts two additional NMS instances.

❖ *To see the current state of NMS, using NMC:*

◆ Type:

```
nmc:/$ show appliance nms pool
```

Editing NMS Properties

You can edit some of the NMS properties to optimize performance of your system.

❖ *To edit NMS properties, using NMV:*

1. Click **Settings > Preferences**.
2. In the **Preferences** panel, click **Server**.
3. In the **Customize NMS Properties** panel, edit the appropriate settings and click **Save**.

❖ *To edit NMS properties, using NMC:*

1. Type:

```
nmc:/$ setup appliance nms property
```

2. Use the arrow keys to select the appropriate setting and make changes as needed.

Table 15-1: NMS Properties

NMS Property	Definition
Auto_scrub_takes_lock	Auto-Scrub service volume-locking policy: Yes No. No (default) — Prevents NMS from locking the volume, when scrubbing. Yes — Locks to prevent other auto-services from running when scrub is in progress
Auto_sync_enforce_ro	Determines whether NexentaStor enforces the destination read-only state. 1 — Prevents any action on the destination dataset, (default). 0 — Means you can destroy or overwrite the destination.
Auto_sync_takes_lock	Auto-Sync service folder-locking policy: Yes No. No (default) — Prevents folder's locking when syncing. Yes — Locks a folder to prevent other auto-services from running when syncing is in progress
Auto_tier_takes_lock	Auto-Tier service folder-locking policy: Yes No. No (default) — Do not lock the folder on local appliance when replicating. Yes — Lock the folder to prevent other auto-services from running on the folder when tiering is in progress.
Automount_timeout	A period in seconds, when a folder remains mounted, when it is not in use. This value applies to NFS and CIFS auto-mounts only. Default = 600 sec.
Autoreserve_space	Allows you to create a reserve volume space in order to solve out-of space issues. NMS creates free-space reserve after volume is 50% full. Default = Yes.
Client_trace_exec_output	Enables tracing of every system exec call on NMS client side.

Table 15-1: NMS Properties

NMS Property	Definition
Client_trace_level	Defines whether to log sysexec perl command, 0 1, 0 — Does not log. 1 — Logs (Default) . <code>sysexec</code> executes various system commands.
Default_ntp_server	Synchronizes time with the specified NTP server.
Disable_plugins_install	Disables plugin installation. Default = No.
Disk_write_cache_mode	Determines how to manage write requests. Before enabling, make sure your SCSI device supports cache settings.
Fast_refresh	Enable to increase speed for refreshing ZFS containers. Default = enabled.
Fast_refresh_min_count	Do not perform if the count of dirty objects is lower than the value. Default = three.
Gid_range_max	The maximum threshold value for GID (Group ID). Default = 60,000.
Gid_range_min	The minimum threshold value for GID (Group ID). Default = 1,000.
Import_caches_timeout	Volume import caches time-out. Default = 3600 seconds.
Index_lang	Sets the language for snapshot suffixes. The default value is english. NexentaStor lists the other available options in NMV.
Indexroot	Specifies a folder to store Indexer's database. Default = <code>syspool/.index</code> .
Internal_broadcast_discovery	Enables NMS to discover all the available hosts for all network interfaces that support broadcast. Yes No. Yes — Enable broadcast discovery. No — Disable broadcast discovery.
Lunsync_on_start	Synchronizes LUNs at NMS startup. If enabled, NMS executes <code>lunsync</code> command at startup, prior to synchronizing the volumes. The option ensures that the kernel discovers certain types of iSCSI and SATA devices properly at boot time. Yes — (default) Enabled. No — Disabled.
net_tcp_naglim_def (NMC only)	Controls Nagle algorithm of coalescing small outgoing TCP packets. Set this value to 1 to disable Nagle; this sends TCP packets without delay. Default = 4095 (NAGLE enabled).
net_tcp_rcv_hiwat (NMC only)	Defines the default receive window size in bytes. The range is from 2048 to 1073741824. Default = 131072.
net_tcp_xmit_hiwat (NMC only)	Defines the default window size in bytes. The range is from 2048 to 1073741824. Default = 131072.
Object_cache_timeout	NMS internal object cache timeout. NMS invalidates its internal per object caches when the specified timeout (in seconds) expires. (Default is 30 sec.)

Table 15-1: NMS Properties

NMS Property	Definition
Rep_rsync_options	Default RSYNC options for tiering service. Change with caution. <ul style="list-style-type: none"> • Delete — Delete files • Exclude — Exclude files from=/var/lib/nza/rsync_excl.txt. • Inplace — Update destination files in the same location. • Ignore — Delete files even if there are I/O errors. • HlptgoD — Instructs rsync to look for hard-linked files when transferring and link them together on the receiving side. Without this option, rsync treats hard-linked files as separate files.
Rep_rsync_version	RSYNC version used for tiering. Default = 3.
Rep_ssh_options	Extra SSH options. Default = empty.
Rsf_config_dist_timeout	The HA cluster configuration update timeout during the HA cluster restart. This property is available if you have the HA cluster plugin installed in your system.
Rsf_config_update_synchronous	HA Cluster configuration update mode. 1 - Strict Synchronous update across the cluster, 0 - Asynchronous update, if synchronous update fails. Default value is 0. Nexenta Systems, Inc. recommends that you do not change the default setting. This property is available if you have the HA cluster plugin installed in your system.
Rsf_default_group_name	Default name of the HA Cluster group. This property is available if you have the HA cluster plugin installed in your system.
Rsf_default_inittimeout	Initial timeout for the HA cluster volume service. Default value is 20 seconds. This property is available if you have the HA cluster plugin installed in your system.
Rsf_default_runtimeout	Timeout for the HA cluster shared service running time. Default value 8 seconds. This property is available if you have the HA cluster plugin installed in your system.

Table 15-1: NMS Properties

NMS Property	Definition
Rsf_log_level	<p>The level of detail of the rsf log file.</p> <p>The options are:</p> <ul style="list-style-type: none"> • 1 — CRITICAL • 2 — CRITICAL and WARNING • 3 — CRITICAL, WARNING, and INFO <p>Default value is 2.</p> <p>This property is available if you have the HA cluster plugin installed in your system.</p>
Sampledoc_maxsize	<p>Maximum indexed document size that is used to generate samples while searching.</p>
Sas2ircu_cache_timeout	<p>Defines the time between requests from NMS to a JBOD chassis using the LSI SAS2 Integrated RAID Configuration Utility (LSI SAS2 IRCU). NMS uses the LSI SAS2 IRCU utility in addition to the SES protocol to get the information about supported JBODs. NMS uses the LSI SAS2 IRCU utility with some HP and Quanta JBODs.</p> <p>Default value is 5 seconds. Nexenta Systems, Inc. recommends that you do not change the default setting.</p>
Saved_configroot	<p>The location to save to and restore from the NexentaStor configuration. Default = <code>.config</code>.</p>
Saved_user_configroot	<p>The location to save to and restore from the NexentaStor configuration. Default = <code>.user.config</code>.</p>
Service_log_keep_days	<p>Determines the period to store log data.</p>
Short_cache_timeout	<p>The minimum time between the requests that NMS sends to a JBOD through the SES protocol. After the <code>short_cache_timeout</code> expires, NexentaStor tries to use the LSI SAS2 IRCU utility to get the information about the JBOD. Only for supported JBODs.</p>
Srvpool_affinity_timeout	<p>Defines the interval for client's requests to be processed with the same NMS instance. Default = 60.</p> <p>For example, you have two NMS instances. A client sends a request to NMS1. If the next client sends a request within 60 seconds after the first request, NMS manages it with the same NMS instance. If the next client sends a request after 60 seconds, NMS transfers it to the next NMS instance.</p>
Srvpool_cnt_initial	<p>The initial number of NMS instances. Default = 2.</p>
Srvpool_cnt_max	<p>The maximum number of NMS instances. Default = 10.</p>
Srvpool_interface_serialization_timeout	<p>The time interval to serialize the API calls through an interface object, provided by a given management server, after any write/delete/create operation occurs on that interface.</p>

Table 15-1: NMS Properties

NMS Property	Definition
Srvpool_port_range_min	Starting value for the range of NMS ports. Default value = 2011.
sys_log_debug (NMC only)	Enables extensive debug-level system logging. 0 — (default), Disables logging. 1 — Enables logging.
sys_snoop_interval (NMC only)	Software watchdog period, in seconds. Default = 50.
sys_snooping (NMC only)	Enable software watchdog. When this is enabled, the appliance reboots if it detects a software error. Default = disabled.
sys_zfs_arc_max (NMC only)	Maximum amount of memory you want the ARC to use, in bytes. ZFS does not necessarily use this much memory. If other applications need memory, the ZFS ARC dynamically shrinks to accommodate the demand. If there is a conflict with applications with a fixed memory demand, you can limit the ARC to reduce or eliminate dynamic memory allocation or deallocation and associated overhead. You can limit the ARC on high-end servers equipped with pluggable CPU/memory boards.
sys_zfs_nocacheflush (NMC only)	Disables device cache and flushes issued by ZFS during write operation processing. This option improves NFS/CIFS performance. Recommended for usage with NVRAM-based cache devices. Caution: Application-level loss of data integrity may occur if appliance is not battery power managed (UPS).
sys_zfs_prefetch_disable (NMC only)	Enables or disables the ZFS intelligent prefetch. When using prefetch, ZFS guesses the next block that the system will read. Generally, ZFS prefetch improves system performance. Default setting is True — enabled. Nexenta Systems, Inc. recommends that you use the default setting.
sys_zfs_resilver_delay (NMC only)	Controls the number of ticks to delay resilvering. Default = 2.
sys_zfs_resilver_min_time (NMC only)	Minimum number of milliseconds to resilver per ZFS transaction group (ZTG).
Timesync_on_start	Synchronizes the time with specified NTP server at NMS startup. Yes No. Yes — (Default) Synchronize. No — Don't synchronize.
Trace_exec_output	Determines whether to trace every system exec call. 0 — Do not trace 1 — (Default) Trace.

Table 15-1: NMS Properties

NMS Property	Definition
Trace_level	The level of system calls tracing. 0 — No tracing 1 — Default 10 — Corresponds to -v command line option 11 — Trace entry/exit 20 — Corresponds to -vv command line option 21 — Backtrace when exception is produced 30 — Corresponds to -vvv command line option.
Uid_admin	User ID for admin. Default = 1000.
Uid_nfs	User ID for nfs User. Default = 6000.
Uid_range_max	Maximum threshold value for UID (User ID). Default = 60000.
Uid_range_min	The minimum threshold value for UID (User ID). Default value = 1000.
Uid_smb	User ID for SMB User. Default = 61,001.
Upgrade_menu_keep	The number of system checkpoints to keep in the GRUB menu. Default = 8.
Upgrade_proxy	A proxy server to download the upgrades in the form of a URL. Default = empty.
Upgrade_rootfs_keep	The number of system checkpoints to store on the system volume. Default = 8.
Upgrade_Type (NMV only)	Remote or local.
volroot	Defines the path where all volumes will be mounted. Default - /volume.

Editing Web UI Properties

You can modify the Web UI settings for NMV.

❖ *To modify the Web UI settings, using NMV:*

1. Click **Settings > Preferences**.
2. Select **Web GUI**.
3. Modify the settings as needed.
4. Click Save.

The following table describes the Web UI settings that you can configure.

Read_access_required	Protect NMV pages that require read access with a password. Default value is <code>No</code> .
Upgrade_check_interval	Interval between upgrade checks. Default value is 0 (disabled).
Appliances_check_interval	Interval between appliance health checks. Default values is 600 seconds.
Pagination_per_page	Number of rows in tables to display per page. Pagination must be enabled. Default value is 15.
Remobj_conn_timeout	Timeout interval for connection between NMV and NMS. Default value is 20 seconds.
Retrieve_treshold	Number of undisplayed items that constitutes a buffer near-empty condition. Default value is 300.
Inc_retrieve	Interval to retrieve buffered charts when the buffer size reaches retrieve threshold. Default value is 60 seconds.
Cpuprof_display_width	Number of <code>dtrace</code> entries that are used to display the CPU status. Default value is 20.
Cpuprof_retrieve_time	Number of <code>dtrace</code> entries to retrieve for the CPU status. Default value is 5.
Tcpstat_display_width	Number of <code>dtrace</code> entries that are used to display the network status. Default value is 20.
Tcpstat_retrieve_time	Number of <code>dtrace</code> entries to retrieve for the network status. Default value is 5.
Iopattern_display_width	Number of <code>dtrace</code> entries that are used to display the disk status. Default value is 20.
Iopattern_retrieve_time	Number of <code>dtrace</code> entries to retrieve for the disk status. Default value is 5.
Analytics_refresh_interval	Interval between analytic data refreshes. Default value is 2000 ms.
Nmv_login_timeout	Disconnect the NMV session after a specified period of time. Default value is 10 minutes.
Ajaxgrid_update_interval	Interval between updates of the AJAX grids in NMV. Default value is 30000 milliseconds.
Ajaxgrid_show_update_link	Determines whether to display the Update link. Default values is Yes .

Managing the Network

This section includes the following topics:

- [About Network Management](#)
- [About Secure Inter-Appliance Access](#)
- [Configuring a Network Interface Card](#)
- [Unconfiguring a Network Interface Card](#)
- [Editing a Network Interface](#)
- [Creating a Link Aggregation](#)
- [Creating an IPMP Network Interface](#)
- [Creating an IP Alias](#)
- [Configuring a VLAN](#)
- [Establishing the SSH-Binding](#)
- [NexentaStor TCP Ports](#)

About Network Management

NexentaStor enables you to perform network management operations, including:

- Managing network interfaces
- Aggregating network interfaces
- Configuring VLANs
- Configuring IP aliases
- Changing TCP ports

Perform the initial network configuration immediately after NexentaStor installation.

About Secure Inter-Appliance Access

NexentaStor provides secure access to other NexentaStor appliances as well as to the administrative management client applications on the network. You can establish the inter-appliance access either using SSH, using the Storage Appliance API (SA-API), or both.

Note:

All management client applications, whether developed internally by Nexenta or by third parties, access other NexentaStor appliances using the SA-API.

In all cases, access to the NexentaStor appliance requires client authentication. NexentaStor supports the following authentication mechanisms:

- Adding the IP address of the client's machine to an IP table
- Using the ssh-binding
- Using the ssh-binding with generated authentication keys

See Also:

- [Establishing the SSH-Binding](#)

Configuring Secure Access Using the IP Table

You can use an IP table to allow root User access to remote NexentaStor appliances without requiring the root User logon and password. This method is not as secure as requiring SSH-key based authentication.

Note: This functionality is only available using NMC.

See Also:

- [NMC Command Line Reference](#)
 - ❖ *To establish the IPv4 inter-appliance authentication, using NMC:*

1. Type:

```
nmc:/$ setup appliance authentication dbus-iptables add allow
```

You can also specify the deny rule for any IP address or subnet mask.

2. Specify the IPv4 subnet mask or IP address of the remote appliance.

Configuring SSH-Key Based Authentication

The most secure method of accessing remote NexentaStor appliances is to provide a unique SSH key on the local appliance and then bind it with the remote appliance. This provides root User access to the remote appliance without requiring the root User password.

Note: This functionality is only available using NMC.

- ❖ *To establish the ssh-key based authentication, using NMC:*

1. Type:

```
nmc:/$ setup appliance authentication keys add
```

2. Enter the **Key ID**.

3. Enter the **Key Value**.

4. After configuring the SSH keys, create the ssh-binding.

See Also:

- [Establishing the SSH-Binding](#)

Configuring a Network Interface Card

You can add additional network interfaces (NIC) to your appliance in order to use them in high availability scenarios, link aggregation, multipath, and so on.

❖ *To add a network interface, using NMV:*

1. Click **Settings > Network**.
2. In the **Network** panel, click **Create**.
3. In the **Create Network Interface** window, select type of the network interface.

The options are:

- Single
 - Aggregated
 - IPMP
4. In the list of devices, select a network interface card.
 5. Optionally, if you select Single or Aggregated interface, specify a link to the interface.
 6. Click **Add Interface**.
 7. Review the results in **Summary Network Settings**.

❖ *To add a network interface, using NMC:*

◆ **Type:**

```
nmc:/$ setup network interface
```

Table 16-1: Network Interface Properties

Network Interface Property	Definition
Interface type	Type of a new network interface. <ul style="list-style-type: none"> • Single — Single physical and logical network interface. Requires one unconfigured NIC. • Aggregated — Technology that combines a few NICs into a single logical device. Requires at least two unconfigured network interfaces. • IPMP, or IP network multipathing — Technology that combines multiple NICs to an IPMP group. The IPMP links can be connected to different switches, which increases availability and performance of the network. You can use IPMP on top of link aggregation.
All available devices	Lists all unconfigured network interfaces available on this appliance.

Table 16-1: Network Interface Properties

Network Interface Property	Definition
Configuration method	A method of NIC configuration. <ul style="list-style-type: none"> • Static — Address assigned by administrator to this particular network interface • Dynamic — Address assigned automatically by a DHCP server when a NIC connects to a network
IP Address	Static unique network address for the network interface. Specify, if you use static configuration method. (Optional parameter.)
Subnet Mask	Static subnet mask for the network interface. Specify, if you use a static configuration method. (Optional parameter.)

Unconfiguring a Network Interface Card

You can unconfigure a NIC interface in order to remove it smoothly from the appliance, or to modify the network.

- ❖ *To unconfigure the network interface, using NMV:*
 1. Click **Settings > Network**.
 2. In the **Summary network settings** panel, in the **Actions** dropdown list, select **Unconfigure** next.
 3. In the confirmation dialog, click **OK**.

- ❖ *To unconfigure the network interface, using NMC:*

- ◆ Type:

```
nmc:/$ setup network interface <interface_name> unconfigure
```

Note: If you unconfigure the primary interface, you lose the connection to NexentaStor over SSH.

Editing a Network Interface

After you create a network interface, you can change its configuration parameters.

- ❖ *To edit the network interface parameters, using NMV:*
 1. Click **Settings > Network**.
 2. In **Summary Network Settings**, click the name of the network interface to edit.

3. In the **Edit Network Interface** panel, perform the required changes and click **Save**.

Tip: You can also unconfigure a network interface on this web-page.
Note that unconfiguring a network interface does not delete VLANs or IP aliases that are configured on the network interface.

4. In the **Configured Network Interfaces** panel, click the interface to edit.

❖ *To edit the network interface, using NMC:*

◆ **Type:**

```
nmc:/$ setup network interface <interface_name>
```

Example:

```
nmc:/$ setup network interface e1000g0
```

Creating a Link Aggregation

NexentaStor fully supports 802.3ad link aggregation. You can use link aggregation to combine multiple physical Ethernet links into one logical link to increase the network performance and to protect the appliance against failures.

Note: You must have at least two unconfigured network interfaces to create link aggregation.

❖ *To create a link aggregation, using NMV:*

1. Click **Settings > Network**.
2. In the **Network** panel, click **Create**.
3. In the **Create Network Interface** window, in the **Interface Type** dropdown list, select **Aggregated**.
4. In the **All Available devices** field, select at least two unconfigured network interfaces.
5. Complete the other fields according to your configuration needs and click **Add Interface**.

The new interface displays in the **Summary Network Settings** panel.

❖ *To create link aggregation, using NMC:*

◆ **Type:**

```
nmc:/$ create network aggregation [-l mode] [-n name] [-P policy] <dev1>  
<dev2> [dev3] ...
```

```
nmc:/$ setup network aggregation create
```

Example:

```
nmc:/$ create network aggregation -l passive e1000g1 e1000g2
```

Tip: Other terms used for link aggregation include Ethernet trunk, NIC teaming, port channel, port teaming, port trunking, link bundling, Ether Channel, Multi-Link Trunking (MLT), NIC bonding, and Network Fault Tolerance (NFT).

Creating an IPMP Network Interface

IPMP, or IP network multipathing provides load balancing, availability, and performance in the systems with multiple interfaces connect to one local area network (LAN).

❖ *To create an IPMP network interface, using NMV:*

1. Click **Settings > Network**.
2. In the **Network** pane, click **Create**.
3. Select the **IPMP** interface type.
4. Select IPMP devices.
5. Click **Add Interface**.

❖ *To create an IPMP network interface, using NMC:*

◆ Type:

```
nmc:/$ setup network ipmp-group create
```

Creating an IP Alias

You can assign more than one IP address to a network interface.

❖ *To create an IP alias, using NMV:*

1. Click **Settings > Network**.
2. In the **Summary network settings** panel, in the **Actions** dropdown list, select **Add IP alias** for the appropriate interface.
3. In the **Create IP alias for Network Interface** window, type an IP alias ID from 1 to 99 and click **Create**.

Tip: The IP alias ID is not the IP address, but an identification of an IP alias.

4. In the **Edit Network Interface** window, choose a configuration method.
 - **Static**
Enter the IP address and subnet mask of an IP alias and click **Save**.
 - **DHCP**
Click **Save**. Network settings are determined automatically by the DHCP server.
 In the **Summary Network Settings**, review the result.

❖ *To create IP alias, using NMC:*

1. Type:

```
nmc:/$ setup network interface
```

2. Select an interface.

Navigate with arrow keys.

3. Select `ipalias`.

4. Type the IP alias ID.

IP alias ID must be an integer value in range from 1 to 99.

Example:

```
IP alias ID : 12
```

System response:

```
e1000g0:12: flags=1001000842<BROADCAST, RUNNING, MULTICAST, IPv4, FIXEDMTU>
mtu 1500 index 2
    inet 0.0.0.0 netmask 0
```

Configuring a VLAN

NexentaStor provides a fully compliant IEEE 802.1Q VLAN implementation. VLAN is a group of hosts with a common set of requirements that communicate as if they were attached to the broadcast domain, regardless of their physical location. You create VLANs to provide the segmentation services traditionally provided by routers in LAN configurations. You can configure VLANs to address the issues such as scalability, security, and network management.

❖ *To configure a VLAN, using NMV:*

1. Click **Settings > Network**.
2. In the **Summary network settings** panel **Actions** dropdown list, select **Add VLAN** for the appropriate interface.
3. In the **Create VLAN for Network Interface** window, type a VLAN ID and click **Create**.
4. In the **Edit Network Interface** window, choose a configuration method.

- **Static**

Enter the IP address and subnet mask of an IP alias and click **Save**

- **DHCP**

Click **Save**. Network settings are determined automatically by the DHCP server.

In the **Summary network settings**, review the result.

❖ *To configure a VLAN, using NMC:*

1. Type:

```
nmc:/$ setup network interface
```

2. Select the name of the interface.

Navigate with arrow keys.

3. Select `vlan`.

4. Type VLAN ID.

VLAN ID must be an integer value in range from 1 to 4095.

Example:

```
VLAN ID : 123
```

System response:

```
e1000g123000: flags=201000842<BROADCAST, RUNNING, MULTICAST, IPv4, CoS> mtu
1500 index 20
        inet 0.0.0.0 netmask 0
        ether 0:c:29:68:14:a8
```

Establishing the SSH-Binding

SSH-binding, aka SSH Public Key Based Authentication, is an easy way to make two NexentaStor appliances communicate through secure connection. When you create an ssh-binding, the public key from the remote NexentaStor appliance is written to the authorized key file on the local appliance.

You need ssh-binding to:

- Manage a remote NexentaStor appliance from the local NexentaStor appliance.
- Create local-to-remote or remote-to-local Auto-Sync replication.
- Create an HA Cluster.

You can create the following types of SSH-bindings:

- **Regular**

The SSH tunnel between two NexentaStor appliances. Use this type of SSH-binding for local-to-remote and remote-to-local Auto-Sync replication, remote appliance management, and other operations.

- **HA-pair**

The SSH tunnel that you create between two NexentaStor appliances for the HA Cluster configuration.

- **VIP**

The SSH tunnel that connects this appliance and a virtual IP address of an HA Cluster group.

❖ *To establish the SSH-binding, using NMV:*

1. Click **Settings > Network**.
2. In the **Network** panel, click **SSH-Bind**.
3. In the **Binding type** field, select the type of binding.

The options are:

- **Regular**
- **HA-pair**
- **VIP**

4. Complete the steps below according to the selected type of the SSH binding.
 - If you selected **Regular**:
 - 1) In the **Remote Server** field, type the IP address of the remote NexentaStor appliance.
 - 2) In the **Remote User** field, type the name of the user with `root` credentials.
 - 3) Type the remote super user password.
 - 4) Optionally, select the Bidirectionally checkbox.
 - If you selected **HA-pair**:
 - 1) In the **Remote Hostname** field, type the hostname of the remote NexentaStor appliance.
 - 2) In the **Remote IP** field, type the IP address of the remote NexentaStor appliance.
 - 3) Optionally, modify the default port number for SSH.
 - 4) In the **Remote User** field, type the name of the user with `root` credentials.
 - 5) Type the remote super user password.
 - 6) Optionally, select the **Bidirectionally** checkbox.
 - If you selected **VIP**: Select the virtual IP address (VIP).

5. Click **Bind**

If you selected **Regular** or **HA-pair**, repeat [Step 1](#) - [Step 5](#) on the remote NexentaStor appliance.

❖ *To establish the SSH-binding, using NMC:*

1. Type:


```
nmc:/$ setup network ssh-bind
```
2. Type the IP address or host name of a remote NexentaStor appliance.
3. Type the remote appliance Super User password.
4. Repeat [Step 1](#) - [Step 3](#) on the remote node

Note:

If ssh-binding fails, you can manually configure the `/etc/hosts/` file, which contains the Internet host table. (In NMC, type `setup appliance hosts` to access the file)

NexentaStor TCP Ports

By default, NexentaStor uses the following TCP ports:

Table 16-2: TCP Port Information

Port	Usage
8457	NMV
2001	NMS
2011, 2012...	Multi-NMS. When multi-NMS is enabled, the port count starts from 2011 and grows respectively.
2002	NMC daemon (NMCd)
2003	Nexenta Management DTrace daemon (NMDTrace)
2049	NFS
4045	NFS
10000	NDMP server
9999	Remote Replication
21/tcp	FTP
22/tcp	SSH
111/tcp	Sun RPC
139/tcp	CIFS (netbios)
445/tcp	CIFS
873/tcp	RSYNC
123/udp	NTP
636/tcp	LDAP
3260/tcp	iSCSI initiator
3205/tcp	iSNS
80/tcp	http
25/tcp	SMTP (fault reporting, tech support requests)

Using NDMP with NexentaStor

This section includes the following topics:

- [About NDMP](#)
- [NDMP Performance Considerations](#)
- [Prerequisites](#)
- [Configuring the NDMP Server](#)
- [Stopping and Restarting the NDMP Server](#)
- [Unconfiguring the NDMP Server](#)
- [Three-Way Copying with NDMPcopy](#)
- [Viewing the NDMP Logs](#)
- [Listing the NDMP Sessions](#)
- [Listing the NDMP Devices](#)
- [Monitoring NDMP Server Performance](#)

About NDMP

Network Data Management Protocol (NDMP) is a networking protocol and an open standard for backing up data in a heterogeneous environment.

Note:

NexentaStor does not provide NDMP management client software. NexentaStor only provides the support of the NDMP protocol and enables NDMP compliant software to copy NexentaStor datasets to tape drives or virtual tape devices.

In order to use NDMP protocol and backup your NexentaStor datasets, you must establish a third-party NDMP server.

The list of available NDMP compliant applications is available at:

<http://www.ndmp.org/products/#backup>

Caution:

Interoperability with 3rd party backup management software is not fully tested. The work is ongoing to test and certify NexentaStor with all major NDMP based backup solutions.

Using NexentaStor with an NDMP compliant backup application allows you to:

- Back up your NexentaStor datasets to tape
- Copy data between two remote servers directly, without having the data traveling through a controlling backup management system.
- Transfer the data on the level of blocks, and on the file level, including the metadata, attributes, ACLs, and so on.

See Also:

- <http://www.ndmp.org/info/overview.shtml>

NDMP Performance Considerations

NDMP performance depends heavily on the average file size. The following list contains an example of NDMP throughputs for different file sizes measured on a Sun Storage 7110 platform:

Table 17-1: Throughput Speeds

File Size	Throughput
8 KB	4.93 MB/sec
64 KB	22.3 MB/sec
128 KB	33.8 MB/sec
1 MB	55.7 MB/sec
1 GB	58.4 MB/sec

During the backup process, the NDMP server uploads the directory structure and file information into the NexentaStor operational memory. Backing up one million files requires approximately 100 MB of memory. If you have a very large number of files to back up, try to back them up in small amounts, such as 10 million files per NDMP session.

During the backup, the NDMP server creates temporary files on disk. You can improve performance by moving the folder for temporary files to the RAM disk.

Note:	<p>The NDMP server requires a minimum of 300 MB to 1 GB of free RAM. If it does not have enough memory, it halts the session and does not write any information to the log file.</p> <p>If you have enough RAM (3GB or more), Nexenta recommends that you move the /ndmp folder for temporary files to the RAM disk, which is mounted as /tmp in NexentaStor. The NDMP server normally uses less than 130 MB for temporary files.</p> <p>If you have several network cards on NexentaStor and want the NDMP server to use only one of them, enter the IP address for the 'mover-nic' option.</p>
--------------	--

Prerequisites

In order to back up NexentaStor datasets to a tape drive or to a virtual tape device, you need:

- An NDMP compliant backup software installed on a non-NexentaStor appliance.
- NexentaStor Enterprise Edition or NexentaStor Trial Version
- A tape drive connected to your NexentaStor appliance or a virtual tape drive. If you connect a tape drive directly to your NexentaStor appliance, NexentaStor can act as a data server.

NexentaStor enables the NDMP management software to copy NexentaStor datasets to the attached tape drive. Your third party client backup software selects to perform all of the management functions, such as setting the replication schedule.

See Also:

- <http://www.ndmp.org/products/#backup>

Configuring the NDMP Server

You must set up an NDMP server on NexentaStor in order to enable the client backup software to transfer NexentaStor datasets. Configuring the NDMP server includes defining the User, password, datapath for temporary files and logs on NexentaStor, along with an optional network interface to use for network traffic.

- The default temporary files directory is `/var/ndmp`. You can change it to `/tmp/ndmp` to improve the performance.
- Temporary files and the log destination files must have at least 300 MB of free space.
- Nexenta recommends that the User name contain at least four characters.
- The password must contain at least eight characters, and can include uppercase symbols and numbers.

Note: Guard the Username and password carefully. Anyone with the Username and password has unrestricted access to all files on all of the NexentaStor volumes.

❖ *To configure the NDMP server, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the **NDMP Server** panel, click **Configure**.
3. In the **NDMP Server: Configure** window, complete the fields to define the service.
4. Click **Apply**.
5. In the **NDMP Server** panel, click **Enable**.

- ❖ *To configure the NDMP server, using NMC:*

1. Type:

```
nmc:/$ setup network service ndmp-server configure
```

2. Follow the prompts to configure the NDMP server.

Stopping and Restarting the NDMP Server

You can stop and restart the NDMP server.

- ❖ *To stop or restart the NDMP server, using NMV:*

1. Click **Settings > Misc. Services**.
2. In the **NDMP Server** panel, click **Disable** (or **Enable**).

- ❖ *To stop or restart the NDMP server, using NMC:*

- ◆ Type:

```
nmc:/$ setup network service ndmp-server disable
or
```

```
nmc:/$ setup network service ndmp-server enable
```

Unconfiguring the NDMP Server

You can remove the NDMP server configuration settings. This returns all NDMP settings to the default values and prohibits access to the NDMP server.

Note: This action is only available using NMC.

- ❖ *To unconfigure the NDMP server, using NMC:*

- ◆ Type:

```
nmc:/$ setup network service ndmp-server unconfigure
```

Three-Way Copying with NDMPcopy

NexentaStor includes an NDMP management utility that can perform NDMP transfers between two NDMP-capable and NDMP-compliant servers. This utility is called `ndmpcopy`.

Note: This action is only available using NMC.

At runtime, `ndmpcopy` establishes connections to the NDMP server processes on the source and destination machines, performs authentication to the servers, and initiates NDMP backup on the source machine and NDMP restore on the destination machine.

Prior to starting NDMP transfer with `ndmpcopy`, prepare the following:

- Source and destination hosts.
- Source and destination folder names.
- Username and password information for both servers.
- ❖ *To start copying, using NMC:*
 1. Type:


```
nmc:/$ ndmpcopy <source_hostname>:/<source_dir>
<destination_hostname>:/<destination_dir>
```
 2. Follow the prompts to specify the User names and passwords for both servers, and to start the data copy.

Viewing the NDMP Logs

The NDMP server uses two log files:

- **system-ndmpd:default.log**
Contains NDMP service lifecycle and service level logging information.
- **ndmplog.0**
Contains NDMP traces and details.
- ❖ *To select the NDMP log files, using NMV:*
 1. Click **Settings > Misc. Services**.
 2. In the **NDMP Server** panel, click **View Log**.
- ❖ *To select the NDMP log files, using NMC:*
 - ◆ Type:


```
nmc:/$ show network service ndmp-server log <logname>
```

Listing the NDMP Sessions

You can select the running NDMP sessions.

Note: This action is only available using NMC.

- ❖ *To list the NDMP sessions, using NMC:*
 - ◆ Type:


```
nmc:/$ show network service ndmp-service sessions
```

Listing the NDMP Devices

You can select any NDMP compliant devices directly attached to NexentaStor or any tape libraries.

Note: This action is only available using NMC.

❖ *To list local NDMP devices, using NMC:*

◆ Type:

```
nmc:/$ show network service ndmp-server devices
```

Monitoring NDMP Server Performance

You can monitor the NDMP server performance.

Note: This action is only available using NMC.

❖ *To monitor the NDMP server performance, using NMC:*

◆ Type:

```
nmc:/$ show network service ndmp-server performance
```

Operations and Fault Management

This section includes the following topics:

- [About Runners](#)
- [About Fault Triggers](#)
- [About Collectors](#)
- [About Reporters](#)
- [About the Indexer](#)
- [Viewing Active Runners](#)
- [Enabling and Disabling Runners and Triggers](#)
- [Changing Runner Parameters](#)

About Runners

NexentaStor runners are services that run in the background and perform essential fault management, performance monitoring, reporting and indexing. Typically, each runner performs a single task.

The runners include:

- **Fault Triggers**
Monitors the system for critical errors.
- **Collectors**
Collects the static information.
- **Reporters**
Sends status reports for the system components.
- **Indexer**
Creates an index for a specified folder to make it available for search.

All runners rely on NMS to provide SA-API. The management console and GUI clients also use this API.

About Fault Triggers

A fault trigger is a runner that performs initial fault management. Every fault trigger uses a certain number of properties and conditions. If any of those conditions occur, a trigger activates and performs the following:

- Sends an e-mail notification with an error report that includes severity, time, scope, suggested troubleshooting action, and often an excerpt of a related log with details.
- Creates an entry in the list of **Faulted Runners** in **Data Management > Runners**.

Note: You must configure the mailer to receive runner notifications.

After a trigger activates and sends the notification, it changes to a maintenance state until you clear all of the errors. The trigger continues to count fault events but does not send additional notifications.

In addition to local fault management, each NexentaStor appliance can monitor other explicitly ssh-bound, or dynamically discovered, appliances on the network using the `hosts-check` trigger.

See Also:

- [Fault ID Reference](#)
- [Setting Up the Mailer](#)

Viewing Fault Triggers

You can view the full list of available fault triggers.

- ❖ *To view all available fault triggers, using NMV:*

- ◆ Click **Data Management > Runners**.

The fault triggers display in the **Registered Fault Triggers** window. You can click a trigger name to display information about the trigger.

- ❖ *To view a specific fault trigger, using NMC:*

- ◆ Type:

```
nmc:/$ show trigger <trigger_type>
```

- ❖ *To view all fault triggers, using NMC:*

- ◆ Type:

```
nmc:/$ show trigger all
```

Viewing Faults

You can display current faults in several ways.

- ❖ *To view all current faults, using NMV:*

- ◆ Click **Data Management > Runners**.

The triggers with faults display in the **Faulted Runners** window. A description of the fault displays.

- ❖ *To view all current faults, using NMC:*

- ◆ Type one of the following commands:

```
nmc:/$ show trigger all faults
nmc:/$ show trigger all-faults all-appliances
```

Clearing Faults

You can clear currently running faults.

❖ *To clear a fault trigger, using NMV:*

1. Click **Data Management > Runners**.

The triggers with faults display in the **Faulted Runners** panel.

2. Select one or more faulted triggers.

3. Click **Clear**.

❖ *To clear a fault trigger, using NMC:*

◆ **Type:**

```
nmc:/$ setup trigger <trigger_type> clear-faults
```

Checking the Fault Management Facility Status

A special trigger checks for NMS and internal system errors. This allows you to monitor the status of the fault management facility itself.

❖ *To check for NMS and internal system errors, using NMV:*

1. Click **Data Management > Runners**.

2. In the **Registered Fault Triggers** panel, click **nms-check**.

❖ *To check for NMS and internal system errors, using NMC:*

◆ **Type:**

```
nmc:/$ show trigger nms-check -v
```

Changing the Content in the Notification Email

If NexentaStor fails and is restarted, you receive an email. The email includes information about the date and time of the reboot, based on how you set verbosity or trace level — the level of detail in the logs. This setting determines how much information the log file or notification contains. You can set up the `trace_level` for almost any trigger.

By default, the last 20 lines of the system log file are included in the email. You can change the amount of data included in the email.

❖ *To change the amount of data included in the restart notification email, using NMV:*

1. Click **Data Management > Runners**.

2. In the **Registered Fault Triggers** panel, click **nms-rebootcheck**.

3. In the **Configurable Properties** window, change the **Tracing Level** as desired.
 4. Click **Save**.
- ❖ *To change the amount of data included in the restart notification email, using NMC:*
1. Type:


```
nmc:/$ setup trigger nms-rebootcheck property
```
 2. Select the `trace_level` option and change it to the appropriate value.

About NM Watchdog

NM Watchdog is a fault trigger that is pre-installed on the NexentaStor appliance. Unlike other runners, you can install and uninstall NM Watchdog using NMC and NMV. The runner monitors the state of NMS and Auto-Sync services. In case of an error, NM Watchdog sends reports by email, writes the failure event in `nms.log`, and creates a notification on the **Status** page in NMV. The NM Watchdog failure report includes actions recommended to fix the failure.

Watchdog runs every 300 seconds.

NM Watchdog verifies the following:

- For NMS:
 - NMS is up and running.

If NMS does not reply to the NM Watchdog requests for 60 seconds for 10 times, NM Watchdog sends a notification that NMS is down. The report includes recommended actions to fix the issue.
- For Auto-Sync:
 - All Auto-Sync services start according to schedule.

If some of the Auto-Sync services did not start on time, NM Watchdog sends a notification with recommended actions to fix the issue. One of the possible reasons why an Auto-Sync service does not start according to schedule is that previous Auto-Sync job has not been finished. Therefore, you may want to adjust replication schedule, so Auto-Sync has sufficient time to complete the replication run before the next replication iteration.
 - Lists of snapshots at source and destination are equal.

With every replication, Auto-Sync first compares the list of snapshots at source and destination. If the lists are not identical, NM Watchdog sends a notification that the Auto-Sync source and destination are not synchronized. Therefore, you may need to verify the state of the source and destination folder to re-enable the Auto-Sync service.

For more information, see NexentaStor Auto-Sync User Guide.

The following text is an example of a report that NM Watchdog sends when NMS is unreachable:

```
Subject: [NMWatchdog Report] host host2

Appliance: host2 (OS version 4.0.1, NMS version 40-0-35)
Machine SIG: 8EFAK8BEA
Timestamp: Mon Mar 3 22:55:28 2014
```


Watcher: NMS

Service name: main

Description: Failed to execute DBus call on NMS: the server does not respond for '60' seconds

NMS process trees:

```
1694 /usr/bin/perl -w /lib/svc/method/nms -d
    2945 <defunct>
    3042 <defunct>
    2907 <defunct>
```

NMS process stack:

```
1694: /usr/bin/perl -w /lib/svc/method/nms -d
    fed67065 pollsys (8047960, 3, 8047a78, 0)
    fed030cf pselect (10, b35d218, fedddf20, fedddf20, 8047a78, 0) + 1bf
    fed033d8 select (10, b35d218, 0, 0, 8047b08, 126e978d) + 8e
    fef15bc8 Perl_pp_sselect (fef7b760, 0, 8047b78, fef76e74, fef7b6ac,
8047e48) + 1a0
    feecbab3 Perl_runops_standard (8047b90, 1, 8047eb8, 0, 8047b90,
feffb0a4) + 27
    fee78d04 perl_run (8061570, 80612d4, 8047dc8, 8051113) + 220
    080511ea main (805128a, feffb0a4, 8047df4, 8050f57, 4, 8047e00) + 11a
    08050f57 _start (4, 8047eb8, 8047ec6, 8047ec9, 8047edd, 0) + 83
```

Suggested possible recovery actions:

- Run 'svcadm clear nms' command
- Reboot the system
- Reboot into a known working system checkpoint

Suggested troubleshooting actions:

- Run 'svcs -vx' command and collect output for further analysis
- Run 'dmesg' command and look for error messages
- Check '/var/log/nms.log' file for error messages
- Check '/var/svc/log/application-nms:default.log' file for error messages

The following text is an example of a report that NM Watchdog sends when Auto-Sync skips a service run according to schedule:

Subject: [NMWatchdog Report] host host3

Appliance: host3 (OS version 4.0.1, NMS version 40-0-35)

Machine SIG: B8A4FJICB

Timestamp: Mon Mar 3 23:06:03 2014

Watcher: AutoSync

Service name: data-0-000

Description: Did not execute on scheduled time: Mon, 03 Mar 2014 23:00:00

MSK: previously scheduled instance of the same service appears to be still running

Suggested possible recovery actions:

- Check the service's schedule and the duration of the replication
- Increase the interval between service runs, and/or tune-up the network configuration to optimize it for replication

```
(hints: multipathing, 10GbE, Jumbo frames, dedicated network)
- Check system date and time
- Check ZFS volumes and folders status
Suggested troubleshooting actions:
- Run 'svcs -vx' command and collect output for further analysis
- Run 'dmesg' command and look for error messages
- Check '/var/log/nms.log' file for error messages
- Check '/var/svc/log/system-filesystem-zfs-auto-sync:data-0-000.log'
file for error messages
```

The following text is an example of a report that NM Watchdog sends when lists of Auto-Sync snapshots at source and destination differ:

```
Subject: [NMWatchdog Report] host host3
X-Mailer: swaks v20130209.0 jetmore.org/john/code/swaks/

Appliance: host3 (OS version 4.0.1, NMS version 40-0-35)
Machine SIG: B8A4FJICB
Timestamp: Tue Mar  4 01:58:50 2014

Watcher: AutoSync

Service name: data-i-000
Description: Destination has missed snapshots: data1/dst/AutoSync-
21_2014-03-04-01-58-03.
Suggested possible recovery actions:
- Check the service's schedule and the duration of the replication
- Increase the interval between service runs, and/or tune-up the network
configuration to optimize it for replication
(hints: multipathing, 10GbE, Jumbo frames, dedicated network)
- Check system date and time
- Check ZFS volumes and folders status
Suggested troubleshooting actions:
- Run 'svcs -vx' command and collect output for further analysis
- Run 'dmesg' command and look for error messages
- Check '/var/log/nms.log' file for error messages
- Check '/var/svc/log/system-filesystem-zfs-auto-sync:data-i-000.log'
file for error messages
```

See Also:

- NexentaStor Auto-Sync User Guide

About Collectors

Collector is a runner that collects and saves the specified statistic in the database.

Table 18-1: Types of Collector and Runners

Runner	Activity
nfs-collector	Collects NFS activity statistical information such as status, state, trace level, period of execution. The runner, <code>nfs-reporter</code> , periodically sends this information in a report.

Table 18-1: Types of Collector and Runners

Runner	Activity
volume-collector	<p>Collects the following volume I/O activity information:</p> <ul style="list-style-type: none"> • reads per second • writes per second • kilobytes read per second • kilobytes written per second • number of waiting transactions • number of active transactions • service time of waiting queue • service time of active queue • queue not empty • transaction in progress • average reads • average writes <p>Periodically <code>volume-reporter</code> sends this information to the configured email address.</p>
smart-collector	<p>Collects SMART disk statistics. <code>smart-collector</code> is a component of the Auto-SMART plugin. The Auto-SMART plugin is included with the NexentaStor software. However, if you uninstall this plugin, the <code>smart-collector</code> is not included in the list of collectors. For more information, see NexentaStor Auto-Smart Plugin User Guide.</p>
network collector	<p>Collects the following network I/O information:</p> <ul style="list-style-type: none"> • input packets • average read in KBytes • input error packets • output packets • average write in KBytes • output error packets • number of collisions

Viewing Collectors

You can view the full list of available collectors.

❖ *To view all available collectors, using NMV:*

◆ Click **Data Management > Runners**.

The collectors display in the **Registered Statistics Collectors** panel. You can click a collector name to display information about the collector.

- ❖ *To view a specific collector, using NMC:*
 - ◆ Type:


```
nmc:/$ show collector <collector_type>
```
- ❖ *To view all available collectors, using NMC:*
 - ◆ Type:


```
nmc:/$ show collector all
```

About Reporters

A reporter is a runner that periodically generates reports about certain parts of the system.

Table 18-2: Runners

Reporter Runner	Activity
config-sd-reporter	Periodically collects support data and saves in the database.
network-reporter	Details information about the size of sent and received packages for every network interface.
nfs-reporter	Sends the NFS statistics collected by the nfs-collector .
services-reporter	Sends the report of every network service for the defined period.
volume-reporter	Sends the volume capacity utilization, history, and statistics report.

Viewing Reporters

You can view the full list of available reporters.

- ❖ *To view all available reporters, using NMC:*
 - ◆ Click **Data Management > Runners**.

The reporters display in the **Registered Reporters** panel. You can click a reporter name to display information about it.
- ❖ *To view a specific reporter, using NMC:*
 - ◆ Type:


```
nmc:/$ show reporter <reporter_type>
```
- ❖ *To view all available reporters, using NMC:*
 - ◆ Type:


```
nmc:/$ show reporter all
```

About the Indexer

The indexer is a runner that allows you to search the existing folders. Creating an indexer means to associate it with a storage folder for the subsequent indexing of this folder and snapshots, and the most recent content. Destroying an indexer means to remove the association with this folder, and, in addition, removing of all the indexed data. Once you index a folder and some/all of its snapshots, you can schedule an indexer to perform the search of the corresponding data. Indexing a large amount of data is a time consuming task, which is why the indexer works offline, according to a defined schedule. However, the search of an indexed folder performs almost immediately.

The Indexer provides the following features:

- Fast search of the indexed folders
- Indexes multiple formats
- NMV and NMC integration
- Smart search, spelling correction suggestions, supports synonyms

The NexentaStor Indexing facility can index the following data:

- AbiWord
- DB2
- HTML
- Microsoft Excel
- Microsoft Powerpoint
- Microsoft Word
- Microsoft Works
- MS SQL
- MySQL
- OpenOffice
- Oracle
- PDF
- Perl
- PHP
- plain text
- POD documentation
- PostgreSQL
- RTF

- SQLite
- StarOffice
- Word Perfect

Viewing Indexers

You can view the full list of available indexers.

- ❖ *To view all available indexers, using NMV:*

- ◆ Click **Data Management > Runners**.

The indexers display in the **Registered Indexers** panel. Click an indexer name to display information about the indexer.

- ❖ *To view a specific indexer, using NMC:*

- ◆ Type:

```
nmc:/$ show indexer <indexer_name>
```

- ❖ *To view all available indexers, using NMC:*

- ◆ Type:

```
nmc:/$ show indexer all
```

Viewing Active Runners

You can display a list of the currently active runners. You can also see the state, status, and schedule of each active runner.

- ❖ *To view the list and status of all current active runners, using NMV:*

1. Click **Data Management > Runners**.
2. To view the status of a specific active runner, click the name of the runner.

- ❖ *To view the list and status of all current active runners, using NMC:*

- ◆ Type:

```
nmc:/$ show appliance runners
```

- ❖ *To view the status of a specific runner, using NMV:*

- ◆ **Type:**

```
show <runner_type> <runner_name>
```

Example:

```
nmc:/$ show reporter volume-reporter
nmc:/$ show trigger nms-check
nmc:/$ show collector volume-collector
nmc:/$ show faults
nmc:/$ show indexer all
```

Some runner types allow the `-v` (verbose) option to display detailed information.

Enabling and Disabling Runners and Triggers

You can enable or disable runners at any time to enhance the performance of your appliance.

- ❖ *To enable or disable all the runners, using NMV:*

1. Click **Data Management > Runners**.
2. In the **Common Runners Operations** panel, click **Enable** or **Disable**.

- ❖ *To enable or disable a specific runner, using NMV:*

1. Click **Data Management > Runners**.
2. In the **Runners** list, click the appropriate runner.
3. Click the **Status** dropdown list, select **enabled** or **disabled**.
4. Click **Save**.

- ❖ *To enable or disable all instances of a specific category of a runner, using NMC:*

- ◆ **Type:**

```
nmc:/$ setup <runner_type> enable|disable
```

Example:

```
nmc:/$ setup collector disable
```

- ❖ *To enable or disable all instances of a specific type of runner, using NMC:*

- ◆ **Type:**

```
nmc:/$ setup <runner_type> <runner_name> enable|disable
```

Example:

```
nmc:/$ setup reporter config-sd-reporter enable
```

Changing Runner Parameters

Every runner has various parameters available for it. You can change a certain parameter to enhance the performance.

❖ *To change the runner parameters, using NMV:*

1. Click **Data Management > Runners**.
2. In the **Runners** list, click the required runner.
3. In the window, change one or more parameters, as available.
4. Click **Save**.

❖ *To change the runner parameters, using NMC:*

◆ **Type:**

```
nmc: /$ setup <runner_type> <runner_name> property
```

Alternatively, you can use the **TAB-TAB** feature to view and select the available runner types, runner names, and properties.

Table 18-3: Runner Properties

Runner Parameter	Definition
Frequency <code>freq_type</code>	Determines the type of the schedule: monthly, weekly, daily, hourly or minute. Depending on this value, you can configure any required schedule.
At day At weekday Period <code>freq_period</code>	Configures this property according to frequency <code>freq_type</code> . Example: <ul style="list-style-type: none"> • If frequency <code>freq_type</code> is hourly, you define the period in hours between service executions. • If frequency <code>freq_type</code> is weekly, you define the day of the week.
At time <code>freq_hour</code> <code>freq_minute</code>	Specifies hour and minute for service execution, when frequency <code>freq_type</code> is daily, weekly or monthly. Specifies minutes for service execution, when frequency <code>freq_type</code> is hourly, daily, weekly or monthly.
Trace Level <code>trace_level</code>	Determines the level of output verbosity.
Output format <code>output_format</code>	Some of the runners specify the type of the report output. Available options: HTML, CSV, ASCII.
Daemon period <code>daemon_period</code>	A period during which a collector takes the samples for the statistics.

Table 18-3: Runner Properties

Runner Parameter	Definition
Upload upload	For some runners, you can save the gathered information to a selected location. The options are: <ul style="list-style-type: none">• 0 — a local drive• 1 — support servers

NMC Command Line Reference

This section includes the following topics:

- [Shell Commands Available in NMC](#)
- [NMC Switch Commands](#)
- [Options Commands](#)
- [Recording for Session](#)
- [NexentaStor Support Request](#)
- [DTrace Framework Commands](#)
- [Query Commands](#)
- [Help Commands](#)
- [Create Commands](#)
- [Destroy Commands](#)
- [Run Commands](#)
- [Share Commands](#)
- [SMTP Commands](#)
- [Setup Commands](#)
- [Show](#)

Shell Commands Available in NMC

NexentaStor passes a number of shell commands through to the underlying bash shell. You can restrict these commands with options. Do not assume they produce the same output as when run on the bash CLI.

```
alias  
awk  
bzip  
cal  
cat  
chmod  
chgrp  
chown  
date  
dd  
dig  
egrep  
fcadm
```

```
fcinfo
find
fmadm
fmdump
grep
gzcat
gzip
history
iostat
ln
lspci
man
mkdir
mkfile
mpathadm
mpstat
ntpdate
ping
prstat
prtconf
prtdiag
ps
quit
rm
rsync
scp
sed
smbadm
sync
tar
tail
top
touch
traceroute
tree
update-pciids
unalias
uptime
vi
vmstat
wc
xargs
zfs
zpool
```

NMC Switch Commands

```
create group basic
```

Options Commands

option *

Recording for Session

record

NexentaStor Support Request

support

DTrace Framework Commands

```
dtrace comstar
dtrace cpu *
dtrace IO*
dtrace locks*
dtrace memory*
dtrace misc*
dtrace network*
dtrace report-all
```

Query Commands

```
usage
auto-service
folder
network-interface
runner
volume
lun
network-service
snapshot
zvol
```

Help Commands

```
help
help advanced-usage
```

```
help commands
help data-replication
help dtrace
help fault-management
help first-steps
help getting-started
help index
help keyword
help more-help
help options
help runners
help tips
```

Create Commands

```
create appliance checkpoint
create appliance user
create appliance usergroup
create folder
create folder show
create group basic
create network aggregation
create script-runner
create ssh binding
create snapshot
create volume
create zvol
```

Destroy Commands

```
destroy appliance checkpoint
destroy appliance user
destroy appliance usergroup
```

Run Commands

```
run *
run diagnostics
run recording
run script-runner
```

Share Commands

```
share * / unshare *
```

```
share folder <vol_name/fol_name>
share folder <vol_name/fol_name> cifs
share folder <vol_name/fol_name> ftp
share folder <vol_name/fol_name> nfs
share folder <vol_name/fol_name> rsync
```

SMTP Commands

```
smtp_addresses
smtp_addresses_faults
smtp_addresses_stats
smtp_auth
smtp_cc
smtp_from
smtp_password
smtp_server
smtp_timeout
smtp_to
smtp_user
```

Setup Commands

All groups of commands relative to setup, except for all plugins, autoservices, indexers, collectors, triggers and reporters.

```
setup appliance *
setup appliance authentication dbus-iptables add
setup appliance authentication dbus-iptables delete
setup appliance authentication dbus-iptables show
setup appliance authentication keys add
setup appliance authentication keys delete
setup appliance authentication keys show
setup appliance checkpoint
setup appliance checkpoint <checkpoint_name> *
setup appliance checkpoint <checkpoint_name> activate
setup appliance checkpoint <checkpoint_name> destroy
setup appliance checkpoint <checkpoint_name> hold
setup appliance checkpoint <checkpoint_name> property *
setup appliance checkpoint <checkpoint_name> property aclinherit
setup appliance checkpoint <checkpoint_name> property aclmode
setup appliance checkpoint <checkpoint_name> property atime
setup appliance checkpoint <checkpoint_name> property canmount
setup appliance checkpoint <checkpoint_name> property checksum
setup appliance checkpoint <checkpoint_name> property compression
setup appliance checkpoint <checkpoint_name> property copies
setup appliance checkpoint <checkpoint_name> property dedup
setup appliance checkpoint <checkpoint_name> property devices
setup appliance checkpoint <checkpoint_name> property exec
setup appliance checkpoint <checkpoint_name> property groupquota
```

```
setup appliance checkpoint <checkpoint_name> property logbias
setup appliance checkpoint <checkpoint_name> property mlslabel
setup appliance checkpoint <checkpoint_name> property nbmand
setup appliance checkpoint <checkpoint_name> property nms:description
setup appliance checkpoint <checkpoint_name> property readonly
setup appliance checkpoint <checkpoint_name> property recordsize
setup appliance checkpoint <checkpoint_name> property refquota
setup appliance checkpoint <checkpoint_name> property refreservation
setup appliance checkpoint <checkpoint_name> property reservation
setup appliance checkpoint <checkpoint_name> property secondarycache
setup appliance checkpoint <checkpoint_name> property setuid
setup appliance checkpoint <checkpoint_name> property show
setup appliance checkpoint <checkpoint_name> property snapdir
setup appliance checkpoint <checkpoint_name> property sync
setup appliance checkpoint <checkpoint_name> property userquota
setup appliance checkpoint <checkpoint_name> property version
setup appliance checkpoint <checkpoint_name> property vscan
setup appliance checkpoint <checkpoint_name> property xattr
setup appliance checkpoint <checkpoint_name> release
setup appliance checkpoint <checkpoint_name> snapshot
setup appliance checkpoint create
setup appliance checkpoint create-from
setup appliance checkpoint restore
setup appliance checkpoint show
setup appliance configuration location
setup appliance configuration restore *
setup appliance configuration restore all
setup appliance configuration restore auto-scrub
setup appliance configuration restore auto-snap
setup appliance configuration restore auto-tier
setup appliance configuration restore basic
setup appliance configuration restore hosts
setup appliance configuration restore interface
setup appliance configuration restore iscsi
setup appliance configuration restore jbods
setup appliance configuration restore mailer
setup appliance configuration restore names
setup appliance configuration restore netsvc
setup appliance configuration restore nmc
setup appliance configuration restore nms
setup appliance configuration restore nmv
setup appliance configuration restore routes
setup appliance configuration restore runners
setup appliance configuration restore shares
setup appliance configuration restore slotmap
setup appliance configuration restore ssh-client
setup appliance configuration restore ssh-server
setup appliance configuration restore users
setup appliance configuration save
setup appliance configuration show
setup appliance domainname
setup appliance domainname show
setup appliance hostname
setup appliance hostname show
```

```
setup appliance hosts
setup appliance init
setup appliance keyboard
setup appliance license
setup appliance license show
setup appliance license -U
setup appliance license -u
setup appliance mailer *
setup appliance mailer show
setup appliance mailer smtp_addresses
setup appliance mailer smtp_addresses_faults
setup appliance mailer smtp_addresses_stats
setup appliance mailer smtp_auth
setup appliance mailer smtp_cc
setup appliance mailer smtp_from
setup appliance mailer smtp_password
setup appliance mailer smtp_server
setup appliance mailer smtp_timeout
setup appliance mailer smtp_to
setup appliance mailer smtp_user
setup appliance mailer verify
setup appliance netmasks
setup appliance nmc edit-settings
setup appliance nmc show-settings
setup appliance nmc-daemon restart
setup appliance nms edit-settings
setup appliance nms pool add
setup appliance nms pool disable
setup appliance nms pool enable
setup appliance nms pool remove
setup appliance nms pool restart
setup appliance nms pool show
setup appliance nms profiler *
setup appliance nms profiler disable
setup appliance nms profiler enable
setup appliance nms profiler reset
setup appliance nms profiler tune
setup appliance nms property *
setup appliance nms property automount_timeout
setup appliance nms property autoreserve_space
setup appliance nms property client_trace_exec_output
setup appliance nms property client_trace_level
setup appliance nms property default_ntp_server
setup appliance nms property disable_plugins_install
setup appliance nms property disk_write_cache_mode
setup appliance nms property fast_refresh
setup appliance nms property fast_refresh_min_count
setup appliance nms property gid_range_max
setup appliance nms property gid_range_min
setup appliance nms property import_caches_timeout
setup appliance nms property index_lang
setup appliance nms property indexroot
setup appliance nms property internal_broadcast_discovery
setup appliance nms property lunsync_on_start
```



```
setup appliance nms property net_tcp_naglim_def
setup appliance nms property net_tcp_recv_hiwat
setup appliance nms property net_tcp_xmit_hiwat
setup appliance nms property object_cache_timeout
setup appliance nms property rep_rsync_options
setup appliance nms property rep_rsync_version
setup appliance nms property rep_ssh_options
setup appliance nms property sampledoc_maxsize
setup appliance nms property saved_configroot
setup appliance nms property service_log_keep_days
setup appliance nms property show
setup appliance nms property srvpool_affinity_timeout
setup appliance nms property srvpool_cnt_initial
setup appliance nms property srvpool_cnt_max
setup appliance nms property srvpool_interface_serialization_timeout
setup appliance nms property srvpool_port_range_min
setup appliance nms property sys_log_debug
setup appliance nms property sys_snoop_interval
setup appliance nms property sys_snooping
setup appliance nms property sys_zfs_arc_max
setup appliance nms property sys_zfs_nocacheflush
setup appliance nms property sys_zfs_resilver_delay
setup appliance nms property sys_zfs_resilver_min_time
setup appliance nms property timesync_on_start
setup appliance nms property trace_exec_output
setup appliance nms property trace_level
setup appliance nms property uid_admin
setup appliance nms property uid_nfs
setup appliance nms property uid_range_max
setup appliance nms property uid_range_min
setup appliance nms property uid_smb
setup appliance nms property upgrade_menu_keep
setup appliance nms property upgrade_proxy
setup appliance nms property upgrade_rootfs_keep
setup appliance nms property volroot
setup appliance nms restart
setup appliance nmv property read_access_required
setup appliance password
setup appliance poweroff
setup appliance reboot
setup appliance register
setup appliance register show
setup appliance remote-access
setup appliance repository
setup appliance repository show
setup appliance swap add
setup appliance swap delete
setup appliance swap show
setup appliance timezone
setup appliance upgrade
setup appliance user create
setup appliance user <username> destroy
setup appliance user <username> property description
setup appliance user <username> property gidNumber
```

```
setup appliance user <username> property homeFolder
setup appliance user <username> property password
setup appliance user <username> property uidNumber
setup appliance user <username> show
setup appliance usergroup <usergroup_name> create
setup appliance usergroup <usergroup_name> create delete-members
setup appliance usergroup <usergroup_name> create destroy
setup appliance usergroup <usergroup_name> create show
setup collector *
setup collector disable
setup collector enable
setup collector network-collector
setup collector nfs-collector
setup collector reset
setup collector run
setup collector show
setup collector volume-collector
setup diagnostics *
setup folder *
setup folder <fol_name> acl
setup folder <fol_name> acl reset
setup folder <fol_name> acl show
setup folder <fol_name> destroy
setup folder <fol_name> destroy-snapshots
setup folder <fol_name> ownership
setup folder <fol_name> ownership show
setup folder <fol_name> property aclinherit
setup folder <fol_name> property atime
setup folder <fol_name> property canmount
setup folder <fol_name> property checksum
setup folder <fol_name> property compression
setup folder <fol_name> property copies
setup folder <fol_name> property dedup
setup folder <fol_name> property devices
setup folder <fol_name> property exec
setup folder <fol_name> property groupquota
setup folder <fol_name> property inherit aclinherit
setup folder <fol_name> property inherit atime
setup folder <fol_name> property inherit checksum
setup folder <fol_name> property inherit compression
setup folder <fol_name> property inherit copies
setup folder <fol_name> property inherit dedup
setup folder <fol_name> property inherit devices
setup folder <fol_name> property inherit exec
setup folder <fol_name> property inherit logbias
setup folder <fol_name> property inherit mslslabel
setup folder <fol_name> property inherit mountpoint
setup folder <fol_name> property inherit nbmand
setup folder <fol_name> property inherit primarycache
setup folder <fol_name> property inherit readonly
setup folder <fol_name> property inherit recordsize
setup folder <fol_name> property inherit secondarycache
setup folder <fol_name> property inherit setuid
setup folder <fol_name> property inherit sharenfs
```

```
setup folder <fol_name> property inherit sharesmb
setup folder <fol_name> property inherit snapdir
setup folder <fol_name> property inherit sync
setup folder <fol_name> property inherit vscan
setup folder <fol_name> property inherit xattr
setup folder <fol_name> property logbias
setup folder <fol_name> property mlslabel
setup folder <fol_name> property nbmand
setup folder <fol_name> property nms:description
setup folder <fol_name> property primarycache
setup folder <fol_name> property quota
setup folder <fol_name> property readonly
setup folder <fol_name> property recordsize
setup folder <fol_name> property refquota
setup folder <fol_name> property refreservation
setup folder <fol_name> property reservation
setup folder <fol_name> property secondarycache
setup folder <fol_name> property setuid
setup folder <fol_name> property show
setup folder <fol_name> property snapdir
setup folder <fol_name> property sync
setup folder <fol_name> property userquota
setup folder <fol_name> property version
setup folder <fol_name> property vscan
setup folder <fol_name> property xattr
setup folder <fol_name> quick-backup
setup folder <fol_name> remount
setup folder <fol_name> share cifs
setup folder <fol_name> share ftp
setup folder <fol_name> share nfs
setup folder <fol_name> share rsync
setup folder <fol_name> snapshot create
setup folder <fol_name> snapshot show
setup folder <fol_name> unshare cifs
setup folder <fol_name> unshare ftp
setup folder <fol_name> unshare nfs
setup folder <fol_name> unshare rsync
setup folder create
setup folder <fol_name> destroy
setup folder show
setup folder version-upgrade
setup group *
setup group basic create
setup group basic show
setup inbox *
setup inbox disable
setup inbox enable
setup inbox mark-all-read
setup inbox show
setup inbox reset
setup inbox set-size
setup indexer create
setup indexer disable
setup indexer enable
```

```
setup indexer reset
setup indexer run
setup indexer show
setup indexer *
setup jbod <jbod_name> model *
setup jbod <jbod_name> model rename
setup jbod <jbod_name> model show
setup jbod rescan
setup lun *
setup lun read_cache disable
setup lun read_cache enable
setup lun show
setup lun slotmap
setup lun slotmap init
setup lun slotmap show
setup lun smart disable
setup lun smart enable
setup lun write_cache disable
setup lun write_cache enable
setup network *
setup network aggregation create
setup network aggregation show
setup network gateway
setup network gateway show
setup network interface *
setup network interface <interface_name> *
setup network interface <interface_name> dhcp
setup network interface <interface_name> ipalias
setup network interface <interface_name> ipalias create
setup network interface <interface_name> linkprops *
setup network interface <interface_name> linkprops adv_autoneg_cap
setup network interface <interface_name> linkprops allowed-dhcp-cids
setup network interface <interface_name> linkprops allowed-ips
setup network interface <interface_name> linkprops autopush
setup network interface <interface_name> linkprops cpus
setup network interface <interface_name> linkprops default_tag
setup network interface <interface_name> linkprops en_1000fdx_cap
setup network interface <interface_name> linkprops en_1000hdx_cap
setup network interface <interface_name> linkprops en_100fdx_cap
setup network interface <interface_name> linkprops en_100hdx_cap
setup network interface <interface_name> linkprops en_10fdx_cap
setup network interface <interface_name> linkprops en_10hdx_cap
setup network interface <interface_name> linkprops flowctrl
setup network interface <interface_name> linkprops forward
setup network interface <interface_name> linkprops learn_decay
setup network interface <interface_name> linkprops learn_limit
setup network interface <interface_name> linkprops maxbw
setup network interface <interface_name> linkprops mtu
setup network interface <interface_name> linkprops pool
setup network interface <interface_name> linkprops priority
setup network interface <interface_name> linkprops protection
setup network interface <interface_name> linkprops rxrings
setup network interface <interface_name> linkprops show
setup network interface <interface_name> linkprops stp
```

```
setup network interface <interface_name> linkprops stp_cost
setup network interface <interface_name> linkprops stp_edge
setup network interface <interface_name> linkprops stp_mcheck
setup network interface <interface_name> linkprops stp_p2p
setup network interface <interface_name> linkprops stp_priority
setup network interface <interface_name> linkprops tagmode
setup network interface <interface_name> linkprops txrings
setup network interface <interface_name> linkprops zone
setup network interface <interface_name> show
setup network interface <interface_name> static
setup network interface <interface_name> unconfigure
setup network interface show
setup network interface vlan
setup network interface vlan show
setup network ipmp-group create
setup network ipmp-group show
setup network nameservers
setup network nameservers show
setup network routes add
setup network routes delete
setup network routes show
setup network service cifs-server *
setup network service cifs-server clear
setup network service cifs-server configure
setup network service cifs-server disable
setup network service cifs-server edit-settings
setup network service cifs-server enable
setup network service cifs-server idmap
setup network service cifs-server join_ads
setup network service cifs-server join_workgroup
setup network service cifs-server log
setup network service cifs-server restart
setup network service cifs-server show
setup network service cifs-server unconfigure
setup network service ftp-server *
setup network service ftp-server clear
setup network service ftp-server confcheck
setup network service ftp-server configure
setup network service ftp-server disable
setup network service ftp-server edit-settings
setup network service ftp-server enable
setup network service ftp-server log
setup network service ftp-server restart
setup network service ftp-server show
setup network service ftp-server unconfigure
setup network service ldap-client *
setup network service ldap-client addcacert
setup network service ldap-client clear
setup network service ldap-client confcheck
setup network service ldap-client configure
setup network service ldap-client delcacert
setup network service ldap-client disable
setup network service ldap-client edit-settings
setup network service ldap-client enable
```

```
setup network service ldap-client log
setup network service ldap-client restart
setup network service ldap-client show
setup network service ldap-client unconfigure
setup network service ndmp-server *
setup network service ndmp-server clear
setup network service ndmp-server confcheck
setup network service ndmp-server configure
setup network service ndmp-server disable
setup network service ndmp-server enable
setup network service ndmp-server log
setup network service ndmp-server restart
setup network service ndmp-server show
setup network service ndmp-server unconfigure
setup network service nfs-client *
setup network service nfs-client clear
setup network service nfs-client confcheck
setup network service nfs-client disable
setup network service nfs-client edit-settings
setup network service nfs-client enable
setup network service nfs-client log
setup network service nfs-client restart
setup network service nfs-client show
setup network service nfs-server *
setup network service nfs-server clear
setup network service nfs-server confcheck
setup network service nfs-server configure
setup network service nfs-server disable
setup network service nfs-server edit-settings
setup network service nfs-server enable
setup network service nfs-server log
setup network service nfs-server restart
setup network service nfs-server show
setup network service nfs-server unconfigure
setup network service nmdtrace-server clear
setup network service nmdtrace-server confcheck
setup network service nmdtrace-server disable
setup network service nmdtrace-server enable
setup network service nmdtrace-server log
setup network service nmdtrace-server restart
setup network service nmdtrace-server show
setup network service nmv-server *
setup network service nmv-server clear
setup network service nmv-server confcheck
setup network service nmv-server disable
setup network service nmv-server edit-settings
setup network service nmv-server enable
setup network service nmv-server log
setup network service nmv-server restart
setup network service nmv-server show
setup network service ntp-client *
setup network service ntp-client clear
setup network service ntp-client confcheck
setup network service ntp-client disable
```

```
setup network service ntp-client edit-settings
setup network service ntp-client enable
setup network service ntp-client log
setup network service ntp-client restart
setup network service ntp-client show
setup network service rr-daemon *
setup network service rr-daemon clear
setup network service rr-daemon confcheck
setup network service rr-daemon configure
setup network service rr-daemon disable
setup network service rr-daemon edit-settings
setup network service rr-daemon enable
setup network service rr-daemon log
setup network service rr-daemon restart
setup network service rr-daemon show
setup network service rr-daemon unconfigure
setup network service rsync-server *
setup network service rsync-server clear
setup network service rsync-server confcheck
setup network service rsync-server configure
setup network service rsync-server disable
setup network service rsync-server edit-settings
setup network service rsync-server enable
setup network service rsync-server log
setup network service rsync-server restart
setup network service rsync-server show
setup network service rsync-server unconfigure
setup network service show
setup network service snmp-agent *
setup network service snmp-agent clear
setup network service snmp-agent confcheck
setup network service snmp-agent configure
setup network service snmp-agent disable
setup network service snmp-agent edit-settings
setup network service snmp-agent enable
setup network service snmp-agent log
setup network service snmp-agent restart
setup network service snmp-agent show
setup network service snmp-agent unconfigure
setup network service ssh-server clear
setup network service ssh-server confcheck
setup network service ssh-server disable
setup network service ssh-server edit-settings
setup network service ssh-server enable
setup network service ssh-server log
setup network service ssh-server restart
setup network service ssh-server show
setup network service syslog-daemon *
setup network service syslog-daemon clear
setup network service syslog-daemon confcheck
setup network service syslog-daemon configure
setup network service syslog-daemon edit-settings
setup network service syslog-daemon enable
setup network service syslog-daemon log
```

```
setup network service syslog-daemon restart
setup network service syslog-daemon show
setup network service syslog-daemon unconfigure
setup network ssh-bind
setup network ssh-bind show
setup network ssh-unbind <host_name>
setup network ssh-unbind <host_name> show
setup network ssl-bind
setup network ssl-bind show
setup network ssl-unbind
setup plugin <plugin_name> *
setup plugin <plugin_name> show
setup plugin <plugin_name> uninstall
setup plugin install
setup plugin show
setup recording *
setup recording show
setup recording start
setup recording stop
setup reporter *
setup reporter config-sd-reporter *
setup reporter config-sd-reporter disable
setup reporter config-sd-reporter enable
setup reporter config-sd-reporter property *
setup reporter config-sd-reporter property freq_day
setup reporter config-sd-reporter property freq_hour
setup reporter config-sd-reporter property freq_minute
setup reporter config-sd-reporter property freq_period
setup reporter config-sd-reporter property freq_type
setup reporter config-sd-reporter property trace_level
setup reporter config-sd-reporter property upload
setup reporter config-sd-reporter reset
setup reporter config-sd-reporter run
setup reporter config-sd-reporter show
setup reporter disable
setup reporter enable
setup reporter network-reporter *
setup reporter network-reporter disable
setup reporter network-reporter enable
setup reporter network-reporter property *
setup reporter network-reporter property freq_day
setup reporter network-reporter property freq_hour
setup reporter network-reporter property freq_minute
setup reporter network-reporter property freq_period
setup reporter network-reporter property freq_type
setup reporter network-reporter property output_format
setup reporter network-reporter property trace_level
setup reporter network-reporter reset
setup reporter network-reporter run
setup reporter network-reporter show
setup reporter nfs-reporter
setup reporter nfs-reporter disable
setup reporter nfs-reporter enable
setup reporter nfs-reporter property *
```



```
setup reporter nfs-reporter property freq_day
setup reporter nfs-reporter property freq_hour
setup reporter nfs-reporter property freq_minute
setup reporter nfs-reporter property freq_period
setup reporter nfs-reporter property freq_type
setup reporter nfs-reporter property output_format
setup reporter nfs-reporter property trace_level
setup reporter nfs-reporter reset
setup reporter nfs-reporter run
setup reporter nfs-reporter show
setup reporter reset
setup reporter run
setup reporter services-reporter *
setup reporter services-reporter disable
setup reporter services-reporter enable
setup reporter services-reporter property *
setup reporter services-reporter property freq_day
setup reporter services-reporter property freq_hour
setup reporter services-reporter property freq_minute
setup reporter services-reporter property freq_period
setup reporter services-reporter property freq_type
setup reporter services-reporter property trace_level
setup reporter services-reporter reset
setup reporter services-reporter run
setup reporter services-reporter show
setup reporter show
setup reporter volume-reporter *
setup reporter volume-reporter disable
setup reporter volume-reporter enable
setup reporter volume-reporter property *
setup reporter volume-reporter property freq_day
setup reporter volume-reporter property freq_hour
setup reporter volume-reporter property freq_minute
setup reporter volume-reporter property freq_period
setup reporter volume-reporter property freq_type
setup reporter volume-reporter property output_format
setup reporter volume-reporter property trace_level
setup reporter volume-reporter reset
setup reporter volume-reporter run
setup reporter volume-reporter show
setup script-runner *
setup script-runner create
setup script-runner disable
setup script-runner enable
setup script-runner reset
setup script-runner run
setup script-runner show
setup snapshot *
setup snapshot <snapshot_name>
setup snapshot <snapshot_name> destroy
setup snapshot <snapshot_name> rename
setup snapshot <snapshot_name> rollback
setup snapshot <snapshot_name> show
setup snapshot create
```

```
setup snapshot show
setup trigger *
setup trigger cpu-utilization-check *
setup trigger cpu-utilization-check clear-faults
setup trigger cpu-utilization-check disable
setup trigger cpu-utilization-check enable
setup trigger cpu-utilization-check property *
setup trigger cpu-utilization-check property cpu_hot
setup trigger cpu-utilization-check property cpu_total_hot
setup trigger cpu-utilization-check property cpu_warm
setup trigger cpu-utilization-check property freq_day
setup trigger cpu-utilization-check property freq_hour
setup trigger cpu-utilization-check property freq_minute
setup trigger cpu-utilization-check property freq_period
setup trigger cpu-utilization-check property freq_type
setup trigger cpu-utilization-check property interval
setup trigger cpu-utilization-check property max_fault_cnt
setup trigger cpu-utilization-check property num_hot_iters
setup trigger cpu-utilization-check property num_procs
setup trigger cpu-utilization-check property num_warm_iters
setup trigger cpu-utilization-check property trace_level
setup trigger cpu-utilization-check reset
setup trigger cpu-utilization-check run
setup trigger cpu-utilization-check show
setup trigger disable
setup trigger disk-check *
setup trigger disk-check clear-faults
setup trigger disk-check disable
setup trigger disk-check enable
setup trigger disk-check property *
setup trigger disk-check property freq_day
setup trigger disk-check property freq_hour
setup trigger disk-check property freq_minute
setup trigger disk-check property freq_period
setup trigger disk-check property freq_type
setup trigger disk-check property ival_check_disk_status
setup trigger disk-check property skip_removable_media
setup trigger disk-check property trace_level
setup trigger disk-check property track_device_not_ready
setup trigger disk-check property track_hard_errors
setup trigger disk-check property track_illegal_requests
setup trigger disk-check property track_media_errors
setup trigger disk-check property track_no_device_errors
setup trigger disk-check property track_prediction_errors
setup trigger disk-check property track_recoverable_errors
setup trigger disk-check property track_soft_errors
setup trigger disk-check property track_transport_errors
setup trigger disk-check reset
setup trigger disk-check run
setup trigger disk-check show
setup trigger enable
setup trigger hosts-check *
setup trigger hosts-check clear-faults
setup trigger hosts-check disable
```

```
setup trigger hosts-check enable
setup trigger hosts-check property *
setup trigger hosts-check property freq_day
setup trigger hosts-check property freq_hour
setup trigger hosts-check property freq_minute
setup trigger hosts-check property freq_period
setup trigger hosts-check property freq_type
setup trigger hosts-check property ival_broadcast_discovery
setup trigger hosts-check property ival_check_ssh_bindings
setup trigger hosts-check property ival_keep_alive_appliances
setup trigger hosts-check property ival_retry_storage_services
setup trigger hosts-check property ival_sync_time
setup trigger hosts-check property nms_down_max_fail
setup trigger hosts-check property timediff_allowed
setup trigger hosts-check property trace_level
setup trigger hosts-check reset
setup trigger hosts-check run
setup trigger hosts-check show
setup trigger memory-check *
setup trigger memory-check clear-faults
setup trigger memory-check disable
setup trigger memory-check enable
setup trigger memory-check property *
setup trigger memory-check property enable_swap_check
setup trigger memory-check property free_ram_notice
setup trigger memory-check property free_swap_critical
setup trigger memory-check property free_swap_notice
setup trigger memory-check property freq_day
setup trigger memory-check property freq_hour
setup trigger memory-check property freq_minute
setup trigger memory-check property freq_period
setup trigger memory-check property freq_type
setup trigger memory-check property interval
setup trigger memory-check property max_fault_cnt
setup trigger memory-check property max_rss_suspect
setup trigger memory-check property num_procs
setup trigger memory-check property paging_critical
setup trigger memory-check property trace_level
setup trigger memory-check property vmstat_aggr_max
setup trigger memory-check reset
setup trigger memory-check run
setup trigger memory-check show
setup trigger nms-check *
setup trigger nms-check clear-faults
setup trigger nms-check disable
setup trigger nms-check enable
setup trigger nms-check property *
setup trigger nms-check property trace_level
setup trigger nms-check reset
setup trigger nms-check run
setup trigger nms-check show
setup trigger nms-fmacheck *
setup trigger nms-fmacheck clear-faults
setup trigger nms-fmacheck disable
```

```
setup trigger nms-fmacheck enable
setup trigger nms-fmacheck property *
setup trigger nms-fmacheck property max_errors_cnt
setup trigger nms-fmacheck property trace_level
setup trigger nms-fmacheck reset
setup trigger nms-fmacheck run
setup trigger nms-fmacheck show
setup trigger nms-rebootcheck *
setup trigger nms-rebootcheck clear-faults
setup trigger nms-rebootcheck disable
setup trigger nms-rebootcheck enable
setup trigger nms-rebootcheck property *
setup trigger nms-rebootcheck property trace_level
setup trigger nms-rebootcheck reset
setup trigger nms-rebootcheck run
setup trigger nms-rebootcheck show
setup trigger nms-zfscheck *
setup trigger nms-zfscheck clear-faults
setup trigger nms-zfscheck disable
setup trigger nms-zfscheck enable
setup trigger nms-zfscheck property *
setup trigger nms-zfscheck property freq_day
setup trigger nms-zfscheck property freq_hour
setup trigger nms-zfscheck property freq_minute
setup trigger nms-zfscheck property freq_period
setup trigger nms-zfscheck property freq_type
setup trigger nms-zfscheck property trace_level
setup trigger nms-zfscheck reset
setup trigger nms-zfscheck run
setup trigger nms-zfscheck show
setup trigger quota-check *
setup trigger quota-check clear-faults
setup trigger quota-check disable
setup trigger quota-check enable
setup trigger quota-check property *
setup trigger quota-check property freq_day
setup trigger quota-check property freq_hour
setup trigger quota-check property freq_minute
setup trigger quota-check property freq_period
setup trigger quota-check property freq_type
setup trigger quota-check property max_fault_cnt
setup trigger quota-check property quota_fault_prc
setup trigger quota-check property quota_info_prc
setup trigger quota-check property trace_level
setup trigger quota-check reset
setup trigger quota-check run
setup trigger quota-check show
setup trigger reset
setup trigger run
setup trigger runners-check *
setup trigger runners-check clear-faults
setup trigger runners-check disable
setup trigger runners-check enable
setup trigger runners-check property *
```

```
setup trigger runners-check property clear_maintenance
setup trigger runners-check property freq_day
setup trigger runners-check property freq_hour
setup trigger runners-check property freq_minute
setup trigger runners-check property freq_period
setup trigger runners-check property freq_type
setup trigger runners-check property max_fault_cnt
setup trigger runners-check property trace_level
setup trigger runners-check reset
setup trigger runners-check run
setup trigger runners-check show
setup trigger services-check *
setup trigger services-check clear-faults
setup trigger services-check disable
setup trigger services-check enable
setup trigger services-check property *
setup trigger services-check property clear_maintenance
setup trigger services-check property freq_day
setup trigger services-check property freq_hour
setup trigger services-check property freq_minute
setup trigger services-check property freq_period
setup trigger services-check property freq_type
setup trigger services-check property max_notify_cnt
setup trigger services-check property trace_level
setup trigger services-check reset
setup trigger services-check run
setup trigger services-check show
setup trigger ses-check *
setup trigger ses-check clear-faults
setup trigger ses-check disable
setup trigger ses-check enable
setup trigger ses-check property *
setup trigger ses-check property freq_day
setup trigger ses-check property freq_hour
setup trigger ses-check property freq_minute
setup trigger ses-check property freq_period
setup trigger ses-check property freq_type
setup trigger ses-check property ival_anti_flapping
setup trigger ses-check property ival_check_ses_sensors
setup trigger ses-check property trace_level
setup trigger ses-check property trace_sensors_unknown
setup trigger ses-check property track_sensors_faults
setup trigger ses-check reset
setup trigger ses-check run
setup trigger ses-check show
setup trigger show
setup trigger volume-check *
setup trigger volume-check clear-faults
setup trigger volume-check disable
setup trigger volume-check enable
setup trigger volume-check property *
setup trigger volume-check property clear_correctable_errors
setup trigger volume-check property enable_ddt_size_check
setup trigger volume-check property enable_mounts_check
```

```

setup trigger volume-check property enable_version_check
setup trigger volume-check property free_space_critical
setup trigger volume-check property free_space_low
setup trigger volume-check property free_space_reserve
setup trigger volume-check property freq_day
setup trigger volume-check property freq_hour
setup trigger volume-check property freq_minute
setup trigger volume-check property freq_period
setup trigger volume-check property freq_type
setup trigger volume-check property ival_check_fma
setup trigger volume-check property ival_check_luns
setup trigger volume-check property ival_check_volumes
setup trigger volume-check property lun_check_model
setup trigger volume-check property lun_check_redundant
setup trigger volume-check property lun_check_size
setup trigger volume-check property max_fault_cnt
setup trigger volume-check property memory_watermark_high
setup trigger volume-check property memory_watermark_low
setup trigger volume-check property trace_level
setup trigger volume-check property use_syspool_check
setup trigger volume-check property volume_version
setup trigger volume-check property volumes_list
setup trigger volume-check reset
setup trigger volume-check run
setup trigger volume-check show
setup usage *
setup volume *
setup volume <vol_name> attach-lun
setup volume <vol_name> clear-errors
setup volume <vol_name> destroy
setup volume <vol_name> destroy-snapshots
setup volume <vol_name> detach-lun
setup volume <vol_name> export
setup volume <vol_name> folder *
setup volume <vol_name> folder <fol_name> *
setup volume <vol_name> folder <fol_name> acl *
setup volume <vol_name> folder <fol_name> acl reset
setup volume <vol_name> folder <fol_name> acl show
setup volume <vol_name> folder <fol_name> destroy
setup volume <vol_name> folder <fol_name> destroy-snapshots
setup volume <vol_name> folder <fol_name> ownership
setup volume <vol_name> folder <fol_name> ownership show
setup volume <vol_name> folder <fol_name> property *
setup volume <vol_name> folder <fol_name> property aclinherit
setup volume <vol_name> folder <fol_name> property aclmode
setup volume <vol_name> folder <fol_name> property atime
setup volume <vol_name> folder <fol_name> property canmount
setup volume <vol_name> folder <fol_name> property checksum
setup volume <vol_name> folder <fol_name> property compression
setup volume <vol_name> folder <fol_name> property copies
setup volume <vol_name> folder <fol_name> property dedup
setup volume <vol_name> folder <fol_name> property devices
setup volume <vol_name> folder <fol_name> property exec
setup volume <vol_name> folder <fol_name> property groupquota

```

```

setup volume <vol_name> folder <fol_name> property inherit
setup volume <vol_name> folder <fol_name> property logbias
setup volume <vol_name> folder <fol_name> property mlslabel
setup volume <vol_name> folder <fol_name> property nbmand
setup volume <vol_name> folder <fol_name> property nms:description
setup volume <vol_name> folder <fol_name> property primarycache
setup volume <vol_name> folder <fol_name> property quota
setup volume <vol_name> folder <fol_name> property readonly
setup volume <vol_name> folder <fol_name> property recordsize
setup volume <vol_name> folder <fol_name> property refquota
setup volume <vol_name> folder <fol_name> property refresreservation
setup volume <vol_name> folder <fol_name> property reservation
setup volume <vol_name> folder <fol_name> property secondarycache
setup volume <vol_name> folder <fol_name> property setuid
setup volume <vol_name> folder <fol_name> property show
setup volume <vol_name> folder <fol_name> property snapdir
setup volume <vol_name> folder <fol_name> property sync
setup volume <vol_name> folder <fol_name> property userquota
setup volume <vol_name> folder <fol_name> property version
setup volume <vol_name> folder <fol_name> property vscan
setup volume <vol_name> folder <fol_name> property xattr
setup volume <vol_name> folder <fol_name> quick-backup
setup volume <vol_name> folder <fol_name> remount
setup volume <vol_name> folder <fol_name> share *
setup volume <vol_name> folder <fol_name> share cifs
setup volume <vol_name> folder <fol_name> share ftp
setup volume <vol_name> folder <fol_name> share nfs
setup volume <vol_name> folder <fol_name> share rsync
setup volume <vol_name> folder <fol_name> show
setup volume <vol_name> folder <fol_name> snapshot
setup volume <vol_name> folder <fol_name> snapshot <snapshot_name> *
setup volume <vol_name> folder <fol_name> snapshot <snapshot_name>
setup volume <vol_name> folder <fol_name> snapshot <snapshot_name> destroy
setup volume <vol_name> folder <fol_name> snapshot <snapshot_name> rename
setup volume <vol_name> folder <fol_name> snapshot <snapshot_name> rollback
setup volume <vol_name> folder <fol_name> snapshot <snapshot_name> show
setup volume <vol_name> folder <fol_name> snapshot create
setup volume <vol_name> folder <fol_name> snapshot show
setup volume <vol_name> folder create
setup volume <vol_name> folder show
setup volume <vol_name> folder version-upgrade
setup volume <vol_name> free-reserve
setup volume <vol_name> grow
setup volume <vol_name> lun read_cache
setup volume <vol_name> lun write_cache
setup volume <vol_name> offline-lun
setup volume <vol_name> online-lun
setup volume <vol_name> property aclinherit
setup volume <vol_name> property atime
setup volume <vol_name> property autoexpand
setup volume <vol_name> property autoreplace
setup volume <vol_name> property canmount
setup volume <vol_name> property checksum
setup volume <vol_name> property compression

```

```
setup volume <vol_name> property copies
setup volume <vol_name> property dedup
setup volume <vol_name> property devices
setup volume <vol_name> property exec
setup volume <vol_name> property groupquota
setup volume <vol_name> property logbias
setup volume <vol_name> property mlslabel
setup volume <vol_name> property nbmand
setup volume <vol_name> property nms:description
setup volume <vol_name> property nms:hold
setup volume <vol_name> property nms:ustatus
setup volume <vol_name> property primarycache
setup volume <vol_name> property quota
setup volume <vol_name> property readonly
setup volume <vol_name> property recordsize
setup volume <vol_name> property refquota
setup volume <vol_name> property refreservation
setup volume <vol_name> property reservation
setup volume <vol_name> property secondarycache
setup volume <vol_name> property setuid
setup volume <vol_name> property show
setup volume <vol_name> property snapdir
setup volume <vol_name> property sync
setup volume <vol_name> property userquota
setup volume <vol_name> property version
setup volume <vol_name> property vscan
setup volume <vol_name> property xattr
setup volume <vol_name> remove-lun
setup volume <vol_name> replace-lun
setup volume <vol_name> show
setup volume <vol_name> spares
setup volume <vol_name> status
setup volume <vol_name> unshare
setup volume <vol_name> version-upgrade
setup volume create
setup volume import
setup volume show
setup zvol *
setup zvol <zvol_name> destroy
setup zvol <zvol_name> property checksum
setup zvol <zvol_name> property compression
setup zvol <zvol_name> property copies
setup zvol <zvol_name> property dedup
setup zvol <zvol_name> property logbias
setup zvol <zvol_name> property nms:description
setup zvol <zvol_name> property primarycache
setup zvol <zvol_name> property readonly
setup zvol <zvol_name> property refreservation
setup zvol <zvol_name> property reservation
setup zvol <zvol_name> property secondarycache
setup zvol <zvol_name> property show
setup zvol <zvol_name> property sync
setup zvol <zvol_name> property volsize
setup zvol <zvol_name> share
```



```
setup zvol <zvol_name> show
setup zvol <zvol_name> snapshot
setup zvol <zvol_name> unshare
setup zvol cache
setup zvol create
setup zvol show
```

Show

```
show all
show appliance authentication dbus-iptables
show appliance authentication keys
show appliance authentication
show appliance auto-services
show appliance checkpoint
show appliance domainname
show appliance dumpdir
show appliance hostname
show appliance hosts
show appliance keyboard
show appliance license
show appliance mailer
show appliance memory
show appliance netmasks
show appliance nmc
show appliance nms
show appliance nmv
show appliance repository
show appliance runners
show appliance saved-configs
show appliance swap
show appliance sysdef
show appliance sysinfo
show appliance syslog
show appliance timezone
show appliance upgrade
show appliance uptime
show appliance user
show appliance usergroup
show appliance uuid
show appliance version
show faults all-appliances
show folder
show group
show inbox
show indexer
show iscsi
show jbod *
show jbod <jbod_name> alerts
show jbod <jbod_name> all
show jbod <jbod_name> blink
```

```
show jbod <jbod_name> fan
show jbod <jbod_name> lun <disk_name> blink
show jbod <jbod_name> lun <disk_name> iostat
show jbod <jbod_name> lun <disk_name> smartstat
show jbod <jbod_name> lun <disk_name> sysinfo
show jbod <jbod_name> lun <disk_name> vtoc
show jbod <jbod_name> model
show jbod <jbod_name> psu
show jbod <jbod_name> sensors
show jbod <jbod_name> slotmap
show jbod alerts
show jbod all
show log *
show log <log_name> log-less
show log <log_name> log-tail
show log syslog dmesg
show log syslog dmesg-tail
show lun *
show lun <lun_name> *
show lun <lun_name> blink
show lun <lun_name> iostat
show lun <lun_name> smartstat
show lun <lun_name> sysinfo
show lun <lun_name> vtoc
show lun cdrom
show lun disk
show lun iostat
show lun slotmap
show lun smartstat
show lun swap
show lun zvol <zvol_name>
show network aggregation *
show network aggregation create
show network aggregation show
show network all-faults
show network devices
show network gateway
show network gateway show
show network interface *
show network interface <interface_name>
show network interface <interface_name> linkprops
show network interface <interface_name> stats
show network interface ipalias
show network interface show
show network interface vlan
show network ipmp-group create
show network ipmp-group show
show network nameservers
show network netgroup
show network routes
show network service
show network ssh-bindings
show network ssl-bindings
show performance
```

```
show performance arc
show performance dtrace
show performance iostat
show performance network
show performance zil
show plugin
show plugin installed
show plugin remotely-available
show recording
show recording all
show script-runner
show share
show snapshot
show usage
show version
show volume
show zvol
```

Fault ID Reference

This section includes the following topics:

- [About Fault ID References](#)
- [Fault Descriptions and Properties](#)

About Fault ID References

The Fault ID References help you identify, track and resolve issues. When an error occurs and a trigger activates, NexentaStor creates a record in its internal fault database in `/var/log/nms.log` and sends an e-mail, if you configured the mailer. (To set up the mailer, see [Setting Up the Mailer](#).)

In NMV, fault notification is displayed under **Fault Management Summary** on the **Status > General** page, as well as on the **Data Management > Runners** page.

Displaying the NMS Log file

The NMS log is the file to which NexentaStor writes records of its activities.

- ❖ *To display the NMS log file, using NMC:*

- ◆ Type:

```
nmc:/$ show appliance nms log-less nms.log
```

- ❖ *To display the most recent entries in the NMS log file, using NMC:*

- ◆ Type:

```
nmc:/$ show appliance nms log-tail nms.log
```

About the NMS fault report

The following text is an example of an `nms.log` fault report:

```

FAULT: *****
FAULT: Appliance   : nexenta1 (OS v3.0.3, NMS v3.0.3)
FAULT: Machine SIG : Z11ZZZZ1Z
FAULT: Primary MAC : 0:15:17:dd:a6:8c
FAULT: Time       : Thu Jun 10 14:15:08 2012
FAULT: Trigger    : services-check
FAULT: Fault Type  : ALARM
FAULT: Fault ID   : 18
FAULT: Fault Count : 1
FAULT: Severity   : NOTICE
FAULT: Description : Service snmp-agent went into maintenance state.
FAULT: Action     : Corrective action applied: service state = online
FAULT: *****

```

The following table describes the NexentaStor fault report.

Table B-1: Fault ID References

Field	Description
Appliance	Appliance host name, OS version, and NMS version
Machine SIG	Machine signature (alphanumeric Machine ID number)
Primary MAC	MAC address used as primary
Time	Exact time that the fault was triggered
Trigger	The trigger or runner that caused this FAULT report (see Operations and Fault Management)
Fault Type	Type of fault: <ul style="list-style-type: none"> • ERROR — Simple fault, such as a command that was executed with an error • ALARM — Serious fault that affects general functionality • FATAL — Severe fault that requires immediate action
Fault ID	The ID of the trigger, indicating a specific reason for the fault.
Fault Count	Number of times the fault occurred
Severity	Level of severity of the fault: <ul style="list-style-type: none"> • INFO — Lowest level, displays the state or action of the current process or service • NOTICE — Second level, describes events that are important, but are not failures • WARNING — Third level, indicates recoverable failure conditions which the system attempts to manage • CRITICAL — Highest level, indicates non-recoverable error requiring administrator action
Description	Short message describing the fault
Action	Suggested action for the administrator, or description of action taken by the system and the result

Fault Descriptions and Properties

This section lists the FAULT_ID values along with the recommended recovery actions.

CPU Utilization Check

Monitors CPU utilization.

Table B-2: CPU Utilization Check

Fault ID Name	ID No.	Description	Recommended Action
Fault ID Exec	1	prstat failed to run Description contains a more detailed message about the error.	Analyze fault report to check syspool current state. Type: <code>nmc:/\$ zpool status</code>
Fault ID Alarm	2	The system is overutilized, remaining idle CPU is __%. The condition persists for more than configured <code>total_time</code> . Utilizes more than indicated <code>CPU_WARM</code> parameter for more than <code>total_time</code> .	Reduce running priority of the process. If this is considered a normal condition, tune up the trigger configuration. Type: <code>nmc:/\$ setup trigger cpu-utilization-check</code> property or see details using: <code>nmc:/\$ show trigger cpu-utilization-check -v</code>

Table B-3: Fault ID Alarm Properties

Property	Default Value	Description
CPU hot	85	Critical threshold for the percentage of CPU that a single process can utilize. Processes using more than this percentage generate a CRITICAL alarm.
CPU total hot	95	Critical threshold for the total percentage of CPU that can be used. When triggered, the runner performs <code>num_hot_iters</code> iterations, each taking an <code>interval</code> of time. It generates a CRITICAL alarm if it exceeds the threshold during all iterations.
CPU warm	70	Warning threshold for the percentage of CPU that a single process can utilize. When triggered, the runner performs <code>num_warm_iters</code> iterations, each taking an <code>interval</code> of time. NexentaStor generates a NOTICE alarm if the threshold exceeds its limit during all iterations.

Hosts Check

Verifies that ssh-bound hosts are kept alive.

Table B-4: Hosts Check Properties

Fault ID Name	ID No.	Description	Recommended Action
Fault Host Pingable Alarm	2	Unable to ping ssh-bound host. This fault occurs when remote host is unreachable for a configured period of time	Verify that remote host is online <code>nmc:/\$ show network ssh-bindings</code>
Fault Time Alarm	3	Time difference number of seconds with remote host exceeded the limit of <code>TIMEDIFF_ALLOWED</code> . Default = 90 seconds	Make sure you set the <code>default_ntp_server</code> property to a known and reachable NTP server hostname or IP address.
Fault Get Time Alarm	4	Unable to retrieve the current time from host	Verify that remote host is reachable and healthy. Type: <code>nmc:/\$ show network ssh-bindings</code>
Fault NMS Down Alarm	5	Nexenta Management Server on the remote host is down	Check <code>/var/log/nms_down.log</code> .
Fault Appliance Pingable Alarm	6	Unable to reach ssh-bounded appliance	Verify that the appliance is online and that its Nexenta Management Server is operational. Type: <code>nmc:/\$ show network ssh-bindings</code>

Table B-5: Fault Appliance Pingable Alarm Properties

Property	Default Value	Description
IVAL Broadcast Discovery	0	Interval of time to periodically run discovery of the appliances on the local network (seconds). This interval is typically on the order of minutes or tens of minutes. Setting this interval to zero effectively disables auto-discovery of appliances.
IVAL Check SSH Bindings	120	Interval of time to periodically validate connectivity to ssh-bound hosts (seconds).

Table B-5: Fault Appliance Pingable Alarm Properties

Property	Default Value	Description
IVAL Keep Alive Appliances	15	Keep-alive appliances interval (seconds). Define the interval as a few seconds. A longer interval can adversely affect group and remote appliance management operations causing User-noticeable delays in response times.
IVAL Retry Storage Services	420	Maximum time interval to retry storage services failed because of the loss of connectivity (seconds). Specifies maximum period of time to validate connectivity to remote hosts associated with (and used by) automated storage services, such as Auto-Tier and Auto-Sync. In most cases the logic attaches to other periodic logic run by this fault trigger, and is retried almost immediately. Note also that loss of connectivity is one of the most common reasons for an automated storage service to fail.
NMS Down Max Fail	3	Maximum number of failures to communicate to local NMS. Number of attempts to try to communicate to local NMS without sending alarm notification.
Timediff Allowed	90	Allowed inter-appliance time difference (seconds). If it exceeds the configured <code>timediff_allowed</code> delta, it triggers an alarm, and the appliances clocks re-synchronize through NTP.

Runners Check

Monitors state and status of the other runners.

In a FAULT report where the trigger is: `runners-check`, disregard the FAULT_ID line with a value higher than 11. Runner-check FAULT ID (except '1' and '11') indicates the indexes of the runners, which are unique for every appliance. If the value is higher than 11, see [Fault Descriptions and Properties](#).

Table B-6: Runners Check

Fault ID Name	ID No.	Description	Recommended Action
Fault ID Eval Error	1	Triggered server error when attempting to check state Description contains more detailed message about the error.	Analyze fault report.

Table B-6: Runners Check

Fault ID Name	ID No.	Description	Recommended Action
Fault ID Maintenance	11	System wasn't able to execute one of the configured runners. Description contains a message about what runner went into maintenance state.	Type: <code>nmc:/\$ show faults</code> Analyze the original cause of the fault. Type: <code>nmc:/\$ setup trigger <trigger></code> <code>clear-faults</code>

Service Check

Monitors state and status of storage services.

In FAULT report where trigger is: `service-check` you can ignore FAULT_ID numbers that are 9 or higher. In this case, the FAULT ID indicates indexes of the runners, which are unique for every appliance. If value is 9 or higher, see FAULT ID SVC MAINTENANCE in the following table.

Table B-7: Storage Service Checks

Fault ID Name	ID No.	Description	Recommended Action
Fault ID Eval Error	1	Triggered server error in attempt to check service state. Description contains more detailed message about the error.	Analyze fault report
Fault ID NMCD Maintenance	2	NMC Daemon went into maintenance state. Description contains last 20 lines of <code>/var/log/nmcd.log</code> .	<code>nmc:/\$ show faults</code> Analyze the log and original fault. Type: <code>nmc:/\$ setup trigger</code> <code>service-check clear-faults</code> Re-enable the runner. Type: <code>nmc:/\$ setup appliance nmc-daemon restart</code>

Table B-7: Storage Service Checks

Fault ID Name	ID No.	Description	Recommended Action
Fault ID NMDTRACE Maintenance	3	Nexenta DTrace Daemon went into maintenance state Description contains last 20 lines of /var/log/nmdtrace.log	Type: <code>nmc:/\$ show faults</code> Analyze the log and original fault. Type: <code>nmc:/\$ setup trigger service-check clear-faults</code> Re-enable the runner. Type: <code>nmc:/\$ setup network service nmdtrace-server clear</code>
Fault ID SVC Maintenance	8	One of the services went into maintenance state. Description contains a message about which service went into maintenance state and 20 last lines of specified service log file.	Type: <code>nmc:/\$ show faults</code> Analyze the log and original fault. Type: <code>nmc:/\$ setup trigger service-check clear-faults</code> Re-enable specified service by clearing maintenance state.

Volume Check

Checks volumes health and capacity, clears correctable device errors, validates mountpoints.

Table B-8: Volume Check

Fault ID Name	ID No.	Description	Recommended Action
Fault ID Eval	1	Server error. Description contains more detailed message about the error.	Analyze fault report.
Fault ID Space	200	Failed to get capacity for volume <vol_name>. Description contains the % of used capacity.	Attach new disk to the volume.
Fault ID Space Perf	300	Volume <vol_name> is low on free space Volume <vol_name>: total free space is only _% of total capacity - performance can degrade.	Attach new disk to the volume.

Table B-8: Volume Check

Fault ID Name	ID No.	Description	Recommended Action
Fault ID Get Status	400	Failed to get volumes status. Description contains more detailed message about the error.	Analyze fault report.
Fault ID Healthy	500	Generates fault report when volume status is not ONLINE and displays the current status. Description contains info about LUNs.	Type: <code>nmc :/\$ show volume status</code> Analyze errors description.
Fault ID Status	600	Status is ONLINE but some errors occurred. One or more devices in volume <code><vol_name></code> experienced errors: read, write, cksum.	System attempts to automatically clear the errors. If attempts failed, then administrative actions required.
Fault ID Mounts	700	Failed to list volume or folder mountpoints or failed to mount. This fault also generates the list of unmounted folders.	Verify mountpoint consistency for volume or folder.
Fault ID Scrub	800		
Fault ID SA API	1000	Failed to execute SA-API trigger method <code>get_faults</code> . Description contains more detailed message about the error.	Analyze fault report.
Fault ID Strict Size	1100	Volume has disks of unequal sizes.	Change volume LUNs.
Fault ID Strict Model	1200	Volume has different disk models.	Change volume LUNs.
Fault ID Strict Redundancy	1300	Volume has non-redundant disk lun.	Add to volume mirror disks.
		Volume groupname configuration consists of too many devices. Recommended maximum number of devices is: RAID-Z1=7, RAID-Z2=9, RAID-Z3=11.	Backup your data and reconfigure the volume.
		Volume has space inefficient groupname configuration. Volume groupname configuration consists of too many devices. Recommended number of devices is 3 or less.	Backup your data and reconfigure the volume.

Table B-8: Volume Check

Fault ID Name	ID No.	Description	Recommended Action
		Volume has space and performance inefficient groupname configuration. Volume groupname configuration consists of too few devices. Recommended minimum number of devices is: RAID-Z1=3, RAID-Z2=4, RAID-Z3=5	Backup your data and reconfigure volume.
Fault ID DDT Size	1400	DDT size surpassed the recommended warning or maximum value. It can cause significant appliance performance decrease.	Add additional RAM or switch off deduplication. To achieve better performance, use SSDs for cache. Extreme values for checking can be manually set. Type: <code>nmc :/\$ setup trigger volume-check property memory_watermark_low high</code> Recommended values are set by default
Fault ID Volume Version (see below)	1500	Compares system version to volume version. Ensures they are the same.	Type: <code>nmc :/\$ nmc :/\$ setup volume <vol_name> version-upgrade</code>

Table B-9: Fault ID Volume Version Properties

Property	Default value	Description
Free Space Critical	5	Critically low free space threshold as percentage of total capacity. Alarm (severity=critical) triggered if the remaining volume capacity is below this threshold.
Free Space Low	20	Low free space threshold as percentage of total volume capacity. Alarm (severity=notice) triggered if the remaining volume capacity is below this threshold.
Memory Watermark Low	50	Specifies recommended warning level % of RAM to be used by DDT. Default = 50%.
Memory Watermark High	80	Specifies recommended maximum % of RAM to be used by DDT. Default = 80%.

Memory Check

Monitors free memory on the appliance.

Table B-10: Memory Check

Fault ID Name	ID No.	Description	Recommended Action
Fault ID Eval	1	Server error. Description contains more detailed message about the error.	Analyze fault report.
Fault ID Alarm Ram	2	Appliance is critically low on free memory. Paging intensity is greater than 10%. Description contains output of command. <code># prstat -s rss -c -n 10 5 1</code>	Add more RAM or increase critical paging threshold by typing: <code>nmc:/\$ setup trigger memory-check property paging_critical</code>
Fault ID Notice Ram	3	Appliance is low on free memory - remains less than 5% of total RAM. Description contains output of command. <code>prstat -s rss -c -n 10 5 1</code>	Add more RAM or increase <code>max_rss_suspect</code> threshold by typing: <code>nmc:/\$ setup trigger memory-check property max_rss_suspect</code>
Fault ID Exec	4	Failed to execute <code>prstat</code> or <code>swap -l</code> . Description contains more detailed message about the error.	Analyze fault report.
Fault ID Alarm Swap	5	Appliance is critically low on swap space - remains less than 10% of total swap space. Description contains output of command. <code># swap -l</code>	Analyze fault report. Increase swap memory might be needed.
Fault ID Notice Swap (see below)	6	Appliance is low on swap space - remains less than 25% of total swap space. Description contains output of command. <code># swap -l</code>	Analyze fault report. Increase swap memory might be needed.

Table B-11: Fault ID Notice Swap Properties

Property	Default value	Description
Free RAM Notice	5	Free memory low threshold (percentage of total RAM). Alarm (severity=notice) triggered if the free memory is below this threshold.
Free Swap Critical	10	Critically low free swap space threshold (percentage of total swap space).
Free Swap Notice	25	Swap space low threshold (percentage of total swap space).
Max RSS Suspect Paging Critical	128	Maximum resident set size (RSS) of any process in the system (megabytes). Alarm (severity=notice) triggered if any process in the system occupies more than the specified RSS.
Paging Critical	10	Critically high paging intensity threshold (percentage of pages loaded from on-disk swap). Alarm (severity=critical) triggered if the intensity of paging to disk (scan rate) is beyond this critical threshold.

License Check

Verifies license info in the system and sends notifications about expiring licenses.

Table B-12: License Check

Fault ID Name	ID No.	Description	Recommended Action
Fault ID License About To Expire	1	Fault starts occurring 7 days before the license expires License-check fault indicates an expiring license, which means that you cannot download upgrades or send support requests. After expiration, your data remains intact and completely accessible through network shares, including NFS and CIFS shares.	You must now stop using the software. Note that you can purchase a new license online or through a third party.
Fault ID Capacity Critical Limit	10	Remaining storage space is less than of maximum allowed by the license.	Order additional capacity online or through partner.
Fault ID Capacity Exceeded	11	Exceeded storage space limit, as per software license.	Order additional capacity online or through partner.
Fault ID Plugins Not Registered	99	One or more plugins are not registered.	Register or re-register plugin.

Table B-12: License Check

Fault ID Name	ID No.	Description	Recommended Action
Fault ID Reregister	100	License expired. Note that the data stored in the appliance remains intact and accessible through network shares, including NFS and CIFS shares.	You must now stop using the software. You can purchase a new license online or through a third party.

NMS Check

Tracks NMS connectivity failures and internal errors. It collects various fault notifications from other runners and sends them to the Administrator.

Table B-13: NMS Check

Fault ID Name	ID No.	Description	Recommended Action
Default NTP Server Fault ID	10	Failed to synchronize time using NTP server.	Check Internet connection.

NMS Autosmartcheck

Checks state of SMART drives.

Table B-14: NMS Auto Smartcheck

Fault ID Name	ID No.	Description	Recommended Action
Fault ID SMART	100	SMART subsystem reported that certain hard disks are generating errors.	Check the disks.

NMS Rebootcheck

Each time NMS starts `nms-rebootcheck`, it checks if the appliance was rebooted.

Table B-15: Checking the Reboot Status

Fault ID Name	ID No.	Description	Recommended Action
Default Reboot Fault ID	1000	Appliance rebooted at specified time and date. There can be many reasons why the appliance rebooted. Fault report contains the last 20 lines of <code>/var/adm/messages</code> .	Analyze fault report.

SES-Check

Verifies the status of JBOD sensors using the SES protocol. In case of a failure, sends a fault report with the name of JBOD, name of the faulted sensor, and its status. Fault ID depends on the type of JBOD. See the JBOD documentation for more details.

Quota-Check

Quota determines the amount of disk space that a user or a group can use in the selected dataset. The quota-check trigger monitors datasets and notifies the NexentaStor Administrator in case a user or a group exceeds the assigned disk space limit.

Table B-16: Quota-Check Fault IDs

Fault ID			
Fault ID Reached Quota	1	Notifies a user that the allocated quota for a dataset has been reached.	Extend the size of the group or user quota for the selected dataset or delete unused data.

Net-Check

Monitors the state of network interfaces and notifies the NexentaStor Administrator about any changes applied to network interfaces. The status of a network interface may change to: `up`, `down`, or `unknown`.

Comstar Trigger

This trigger runs, if the cluster update fails:

- Failed to revert the nodes to a previous state.
- Operation succeeded on the local node, but failed on the remote node.
- Failed to update all nodes.
- Does not have any fault IDs or properties.

NMS ZFS Check

Monitors the changes in ZFS and updates the cache files, if a volume is created or removed. This is an internal trigger. Used by the disks managing the module. It does not have any fault IDs or properties.

NMS FMA Check

Delivers FMA events through configured transports. Parses log files associated with Illumos Fault Manager and displays errors from them.

FMA report notifies you about faults in your system.

NMS FMA Check receives error reports from the following diagnose engines:

- **Slow I/O Diagnose Engine (DE)** or `slow-io-de`

The Slow I/O Diagnose Engine monitors the I/O time of disk drives. Typically, for an HDD the standard I/O response time is approximately 2 ms. The Slow I/O Diagnose Engine notifies you about a potential hard drive issue if the I/O response time is longer than 2 seconds in 10 I/O requests during a period of 10 minutes. `nms-fma-check` sends a fault report with the description of the problem and actions taken to eliminate the problem.

The following issues in your system may trigger the Slow I/O event notification:

- Damaged cable or expander
- Disk failure
- External vibration

The FMA service may perform the following actions:

- Mark disk as faulted and taken out of service

You may get this message if you have a redundant pool configuration with at least one disk in operation.

- Mark disk as faulted and leave it in operation.

You get this message if you have no healthy disks in a pool. Therefore, data remains available to network clients.

- **The Disk-Sense Diagnose Engine (DE)** or `disk-sense-de`

Hard drives may return sense codes in case of a potential disk drive issue. When the error count exceeds the threshold of 10 device errors in 10 minutes, the FMA service tries to retire the disk. The FMA service also generates a report that contains a key code qualifier that describes the error event.

For more information about Sense Code Qualifiers, see http://en.wikipedia.org/wiki/Key_Code_Qualifier.

Note: Nexenta Systems, Inc. recommends that you replace any disk drive that was diagnosed as faulted.

Disk Check

Monitors the state of the disk. Does not have any fault IDs. Used by the disks managing the module. Disabled by default.

Table B-17: Disk Check Properties

Fault ID Name	Description
IVAL Check Disk Status	Interval (seconds) to periodically check physical status of disks
Track Soft Errors	Tracks soft errors. A soft error is a general concept which includes other types of repairable errors, such as recoverable errors, illegal request errors, transport error, and so on.
Track Hard Errors	Tracks hard errors. A hard error is a general definition of the error that cannot be easily recovered and requires some additional administrative actions. Such errors as: media, no device, device not ready errors.
Track Transport Errors	Tracks transport errors. The error can happen because of the inability to transport the command.
Track Recoverable Errors	Tracks the errors that can be recovered. Recoverable error means that the command was successfully transported, but some additional actions are required.
Track Prediction Errors	Tracks predictive failure analysis errors. This property informs you that currently the disk does not have any errors, but it might have them soon.
Track No Device Errors	Tracks <code>No Device</code> errors. Device cannot be located.
Track Media Errors	Tracks media errors.
Track Illegal Requests	Tracks illegal requests. Illegal requests can be issued by improper commands or a software bug.
Track Device Not Ready	Tracks <code>Device Not Ready</code> errors. The media device (SCSI or FC) has gone offline or is dead. Check the cable connections.

Auto-Sync SVC Trigger

Auto-Sync creates a trigger for every service. The trigger monitors work of the Auto-Sync service and notifies NexentaStor Administrator about failures.

For every error, Auto-Sync creates a fault report that contains the following information:

- **Error source**

Error source may include: `autosync`, `rrdaemon`, or `rrmgr`.

- **Error type**

Error type depends on the error source.

- If the error source is `autosync`, the error type is undefined.
- If the error source is `rrdaemon` or `rrmgr`, the error type may be defined as:
 - `command`
 - `stream`
 - `transport`

- **Fault ID**

Fault ID is a number that defines the error. See [Auto-Sync Fault IDs](#). If the fault ID number is in range 1 — 150, the fault is related to a system call error. If the fault ID number is in range 2000 — 2070 and the source is `rrdaemon` or `rrmgr`, then the fault is related to `libzfs`. For more information, see the ZFS documentation.

- **Message**

An error message provides details about the fault.

Table B-18: Auto-Sync Fault IDs

Fault ID Name	Fault ID	Description	Recommended Action
Replication Error	255	An undefined error occurred during the replication.	Verify that source and destination appliances are available. Verify that the Auto-Sync service instance is enabled. Enable if required.
EIO	5	Input/Output error. Fault report contains more information about the error.	This error indicates the issue on a disk level.
EFAULT	14	A system operation failed.	Fault report contains more information about the error.
EINVAL	22	Invalid argument.	Verify that parameters specified for the Auto-Sync service are valid.
ENOSPC	28	No space left on device	Free additional space at destination.
ENODATA	61	The dataset that Auto-Sync needs to replicate is not found.	Verify that the source dataset is available.

Table B-18: Auto-Sync Fault IDs

Fault ID Name	Fault ID	Description	Recommended Action
ETIMEDOUT	110	Connection time out	Verify network connectivity between two appliances.

See Also:

- http://www.numi.fnal.gov/offline_software/srt_public_context/WebDocs/Errors/unix_system_errors.html

Index

- A**
- ACL 68
 - managing 83
 - permissions 84
 - active directory 97
 - adding NexentaStor 99, 111
 - integration 111
 - joining 101
 - prerequisites 98
 - verifying settings 98
 - active runners
 - viewing 213
 - adding
 - spare device
 - existing volume 54
 - volume creation 53
 - analytics settings
 - changing 163
 - anonymous
 - user access
 - CIFS 78
 - appliance
 - powering off 17
 - rebooting 17
 - attaching
 - disk to volume 36
 - authenticated access
 - workgroup mode 79
 - authentication
 - chap 133
 - domain mode 101
 - workgroup mode 94
 - auto-scrub
 - creating 47
 - auto-service
 - creating 152
 - deleting 155
 - disabling 154
 - editing 152
 - enabling 154
 - starting 154
 - viewing 152
 - auto-services 147
 - 151
 - auto-scrub 151
 - auto-smart 152
 - auto-snap 151
 - replication 148
 - auto-sync 148, 150
 - auto-tier 148
 - properties 149
- C**
- CHAP
 - authentication 133
 - bidirectional 133
 - chart profiles 163
 - checkpoint
 - changing 143
 - creating 141
 - deleting 143
 - rolling back 142
 - saving 143
 - viewing 142
 - checkpoints 136, 141
 - CIFS 93
 - anonymous
 - access 78
 - authentication
 - domain 101
 - domain mode 94
 - workgroup
 - authentication 94
 - workgroup mode 93
 - CIFS server
 - configuring 75
 - options 76
 - clearing faults 206
 - collectors 209
 - viewing 210
 - commands
 - displaying 11
 - at prompt 12
 - shortcuts 12
 - NMC
 - create 220, 221
 - destroy 220
 - dtrace 219
 - help 219
 - options 219
 - query 219
 - record 219
 - run 220
 - share 220
 - show 240
 - SMTP 221
 - support 219
 - switch 218, 220, 221
- comparing
 - auto-tier | auto-sync 148
- components
 - NexentaStor 2
- configurations
 - restoring 144
 - saving 144
- configuring
 - CIFS server 75
 - FTP server 79
 - LDAP server 120
 - NFS server 73
 - NXS 103
 - Windows
 - for AD 101
- controlling access 68
- creating
 - auto-scrub 47
 - folders 69
 - local groups 118
- custom chart
 - removing 163
- D**
- data volumes
 - viewing
 - status 35
 - default users 114
 - deleting
 - folders 90
 - volume 46
 - detaching
 - disk

- mirror 42
- device
 - replacing
 - pool 44
 - setting
 - offline 45
 - online 45
- device types 6
- disk
 - detaching
 - mirror 42
- disk management 51
- displaying
 - properties
 - shortcut 13
 - volume properties 37
- domain mode 94
 - authentication 101
- E**
- editing
 - folder properties 86
 - volume properties 38
- email
 - notification 206
- emergency reserve
 - disabling 49
 - enabling 49
- existing volume
 - spare device
 - adding 54
- expert mode
 - using 11
- exporting
 - volume 47
- F**
- fault ID
 - alarm
 - properties 245
 - alarm properties 245, 246
 - appliance alarm
 - properties 246
 - comstar trigger 256
 - CPU utilization check 245
 - disk check 257
 - hosts check 246
 - properties 246
 - license check 253
 - memory check 252
 - NMS autosmartcheck 254
 - NMS check 254
 - NMS FMA check 256
 - NMS rebootcheck 255
 - NMS ZFS check 256
 - notice swap
 - properties 253
 - service check 248
 - volume check 249
 - volume properties 251
- fault management 206
- fault triggers 204
 - viewing 205
- faults
 - clearing 206
 - viewing 205
- folders
 - creating 69
 - deleting 90
 - editing properties 86
 - indexing 90
 - management 67
 - searching 91
 - shared
 - controlling access 68
 - sharing 68
 - CIFS 75, 76
 - ftp 79, 80
 - NFS 74
 - RSYNC 81, 82
 - sharing options 75
 - tasks 85
 - unsharing 85
 - viewing status 71
- FTP
 - sharing folders 79, 80
- FTP server
 - configuring 79
- G**
- growing
 - volume 36
- guide
 - audience xv
- I**
- I/O performance
 - viewing 163
- ID mapping 104
- importing
 - volume 48
- indexer 212
- indexers
 - viewing 213
- indexing
 - folders 90
- initiator group
 - initiating 135
- instance 147
- IP table 189
- iSCSI target
 - creating 127
- L**
- LDAP
 - example 122
 - server
 - configuring 120
 - servers
 - supported 120
 - users 119
- local
 - user groups 117
- local groups
 - creating 118
 - managing 119
- logical dataset 147
- LUN mapping 131
 - creating 132
- M**
- mailer
 - setting up 19
- managing
 - ACL 83
 - folders 67
 - local groups 119
- mode
 - domain 93, 94
 - workgroup 93
- mounting
 - shared folders
 - Windows 83
- Multi-NMS 180
- multiple options
 - selecting 13
- N**
- NDMP 198
 - configuring 200
 - devices

- listing 203
- logs 202
- monitoring performance 203
- performance considerations 199
- prerequisites 200
- restarting server 201
- sessions
 - listing 202
 - stopping server 201
 - unconfiguring 201
- NDMPcopy 201
- network management 188
- NexentaStor 1
- NexentaStor components 2
- NFS
 - sharing folders 72
- NFS server
 - configuring 73
 - tunable options 73
- NIC
 - configuring 190
 - unconfiguring 191
- NMC 2, 9
 - commands
 - reference 14
 - create 220, 221
 - destroy 220
 - dtrace 219
 - expert mode
 - using 11
 - Help
 - displaying 14
 - help 219
 - logging in 9
 - ssh 9
 - logging out
 - local 10
 - remote 10
 - navigating 11
 - options 219
 - query 219
 - record 219
 - run 220
 - setup 221
 - share 220
 - shell commands 217
 - show 240
- SMTP 221
 - support 219
 - switch commands 218
- NMS 2, 180
 - editing properties 181
 - Multi-NMS 180
 - restarting 22
- NMV 2
 - displaying URL 10
 - exiting 9
 - getting help 9
 - logging in 8
 - logging out 8
 - refreshing 9
 - starting 7
 - user permissions 115
 - using 7
- NXS
 - configuration settings
 - modifying 18
 - configuring 103
 - services
 - restarting 21
- NXS environment
 - planning 3
- O**
 - options
 - displaying 11
 - shortcut 13
- P**
 - performance benchmarks 164
 - planning
 - NXS environment 3
 - plugins 2
- Q**
 - quota
 - setting 118
- R**
 - redundancy groups 5
 - remote initiator 134
 - adding 134
 - removing device
 - volume 43
 - replacing
 - device
 - pool 44
 - reporters 211
 - viewing 211
- RSYNC
 - sharing
 - folders 81, 82
 - tunable options 81
 - runner parameters
 - changing 215
 - runners 204
 - runners and triggers
 - disabling 214
 - enabling 214
- S**
 - scripting 17
 - SCSI target 126
 - searching
 - folders 91
 - secure authentication
 - chap 133
 - setting
 - device
 - offline 45
 - online 45
 - group quota 118
 - group size limit 118
 - setting permissions 100
 - settings
 - changing manually 18
 - changing with wizards 18
 - shared folders
 - controlling access 68
 - sharing
 - folders
 - CIFS 75, 76
 - NFS 74
 - RSYNC 81, 82
 - sharing folders 68
 - FTP 79, 80
 - NFS 72
 - size limit
 - setting 118
 - snapshots 136
 - cloning 138
 - creating 137
 - deleting 140
 - recovering files 139
 - renaming 141
 - service
 - setting up 136
 - using 137

- viewing data 138
- spare device
 - adding
 - existing volume 54
 - volume creation 53

- SSH
 - configuring 189

- statistics
 - available 159
- NMC
 - CPU 161
 - IO 161
 - locks 162
 - memory 162
 - miscellaneous 162
 - network 162
- performance 157
- viewing 157

- storage reqs 3
- support
 - contact xv
- supported
 - LDAP
 - servers 120

T

- target group
 - creating 128
- target portal group
 - creating 127
- troubleshooting 109
 - general 110
- tunable options
 - NFS server 73

U

- unsharing
 - folders 85
- user accounts
 - default 15
 - default pwd
 - changing 16
 - passwords
 - default 16
 - user names
 - default 16
- user groups
 - local 117
- user permissions 114
- users

- LDAP 119
- using
 - expert mode 11
- using volumes
 - main tasks 29

V

- viewing
 - active runners 213
 - collectors 210
 - fault triggers 205
 - faults 205
 - folder status 71
 - indexers 213
 - reporters 211
 - status
 - data volumes 35

- volume
 - additional tasks 37
 - attaching disk 36
 - creating 29
 - deleting 46
 - exporting 47
 - growing 36
 - importing 48
 - properties 86
 - removing device 43
- volume creation
 - spare device
 - adding 53
- volume management 28
- volume properties
 - displaying 37
 - editing 38

W

- Windows
 - active directory 97
 - configuring
 - for AD 101
 - mounting
 - shared folders 83
- workgroup
 - authentication 94
- workgroup mode 93
 - authenticated access 79

Z

- ZVOL 128
 - creating 129
 - destroying 131

- properties
 - changing 130
 - viewing 130

Global Headquarters

451 El Camino Real, Suite 201
Santa Clara, CA 95050
USA

3000-nxs-4.0.4 FP1-000057-A