



Global Leader in Software-Defined Storage.

NexentaStor 5.0 CLI Configuration QuickStart

Date: October, 2016

Part Number: 3000-nxs-cliconfig-5.0-000092-A

Copyright © 2016 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, NexentaFusion, 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.

Contents

Preface	vi
1 Initial Configuration	1
Introduction	1
Using the CLI	1
Post-Install Checks	2
Changing Management Address	3
Managing Users	4
Configuring SMTP Email Service	5
What Comes Next?	7
2 Configuring NexentaStor 5.0 Appliances	8
Summary of Storage and Cluster Configuration	8
Verifying Enclosure and Disk Information	11
Managing Network	13
Verifying the Available Network Interface	13
Assigning Addresses to Network Interfaces and Servers	13
Creating a Network Route	14
Aggregating NICs	15
Using VLANs	15
Configuring HA Cluster	16
Configuring a Pool	19
Creating a Pool	19
Adding Spares, Cache Devices, and Log Disks	21
Smart-Sparing	22
Maintaining Pools (Scrub and Trim/Unmap)	22
Configuring HA Service for the Shared Pool	24
Creating and Sharing a File System	27
Creating a File System	27
Sharing a File System via SMB	28
Sharing a File System via NFS	31

Creating and Sharing Volumes	36
Creating Volume Groups and Volumes	36
Sharing iSCSI Volumes	37
Sharing to Any Host on the Network	38
Sharing to a Specific Set of iSCSI Initiators	39
Preparing iSCSI LUNs for Failover	39
Sharing FC Volumes	40
Sharing to a Specific Set of FC Initiators	40
Preparing FC LUNs for Failover	41
What Comes Next?	42
3 Data Protection	43
Protecting Data	43
Configuring Snapshots and Snapshot Schedules	44
Creating and Deleting Local Snapshots	44
Creating a Local Scheduled Snapshot Service	45
Managing Snapshots	46
Configuring High Performance Replication	48
About HPR	48
Definition of Terms	49
Preparing for HPR	49
Licensing Requirements for the Replication Service	52
Managing the HPR Password	52
Configuring HPR Address to Send or Receive Replication Traffic	53
Configuring a New Replication Service	54
Managing the Created Replication Service	57
Scheduling Replication Services	58
Configuring NDMP for Backups	60
Prerequisites	60
Configuring and Managing NDMP	60
Troubleshooting	60
Using Virus Scan Engines	62
About NexentaStor vscan	62
Prerequisites	62

Managing Virus Scan Services on a NexentaStor Appliance	62
Checking vscan Related Messages in the NexentaStor Syslog	63
What Comes Next?	64
4 Fault Management	65
Viewing NexentaStor Logs	65
Troubleshooting Data Devices	67
Alerts and Events	69
Diagnose Unsuccessful HA Service Start	71
Bundle Services	72
Creating a Support Bundle	72
Verifying the Bundle	72
Uploading a Support Bundle	72
Upload Status	72
What Comes Next?	73
5 Advanced Configuration	74
Configuring Write-Back Cache	74
Sharing NFS and SMB Concurrently on a File System	77
Creating Name-Based Rules	77
Using IPMP	78
Additional Resources	79

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-cliconfig-5.0-000092-A	October, 2016	GA

Contacting Support

Send your support questions and requests to support@nexenta.com.

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.

Initial Configuration

This chapter covers the following topics:

- [Introduction](#)
- [Using the CLI](#)
- [Post-Install Checks](#)
- [Changing Management Address](#)
- [Managing Users](#)
- [Configuring SMTP Email Service](#)
- [What Comes Next?](#)

Introduction

This document is intended to be read after installing NexentaStor appliance in your environment. This guide demonstrates the basic steps and commands to configure and manage NexentaStor 5.0 appliances using the CLI. We recommend that you use this document in conjunction with the following guides:

- [NexentaStor 5.0 CLI Reference Guide](#)
- [NexentaStor 5.0 HA Admin Guide](#)

Using the CLI

For additional help options when using the CLI commands, use the following:

<code>Ctrl+C</code>	Returns you to the prompt.
<code>CLI@nexenta> help</code>	Lists available CLI commands and UNIX-like utilities.
<code>CLI@nexenta> man cli</code>	Provides an overview of the CLI commands and the general options and output flags.
<code>CLI@nexenta> man <command></code>	Displays the man page for a specific command. Use the Spacebar or arrows to move through the man page. Press <code>q</code> to return to the prompt.
<code>CLI@nexenta> <command></code>	Displays available subcommands for the specified command.

CLI@nexenta> <command> Provides usage information for the specified
 <subcommand> --help subcommand

Post-Install Checks

Before you proceed with configuring pools, file systems, or volumes, review your baseline settings after installation from the CLI using the following commands.

Task	Related CLI Command
Verify software version installed	CLI@nexenta> software version
Verify the installed software and the boot environment	CLI@nexenta> software list
List available licenses	CLI@nexenta> license show To use the High Availability (HA), continuous replication, fibre channel management, and All-flash optional features, license tokens for these must be activated and listed in the <code>features</code> row as shown in the following sample output. If these features are not enabled, contact sales@nexenta.com to obtain the feature license. Sample output: <pre> PROPERTY VALUE guid 534d4349-0002-6190-2500-6190250015b0 valid yes status ok type ENTERPRISE-TRIAL (Nexenta Internal) product NexentaStor version 5.x licensee Nexenta-abcd@nexenta.com serial SR-DEV-NS-201616578 features allFlash, continuousReplication, scheduledReplication, fibrechanne issued Wed Sep 21 17:00:00 2016 expires Thu Dec 29 16:00:00 2016 capacity no limit maintenance Wed Sep 21 17:00:00 2016 - Thu Dec 29 16:00:00 2016 (valid) </pre>
See your system's baseline state	CLI@nexenta> system status
Check that the system services of interest are enabled	CLI@nexenta> svc list

Check that NTP service is enabled and settings are correct	<pre> 1. Set up NTP CLI@nexenta> svc set servers=pool.ntp.org ntp 2. Check if ntp is online CLI@nexenta> svc list ntp NAME DESCRIPTION STATE ntp NTP client online 3. Validate that the ntp server is set. CLI@nexenta> config list system.date CLI@nexenta> svc get ntp Name Value servers ['pool.ntp.org'] </pre>
List enclosures	<pre>CLI@nexenta> enclosure list</pre>
List available disks	<pre>CLI@nexenta> disk list</pre>
List available NICs	<pre>CLI@nexenta> inventory nic</pre>
List available links	<pre>CLI@nexenta> link list</pre>
List current routes	<pre>CLI@nexenta> route list</pre>
List configured IP addresses	<pre>CLI@nexenta> ip list</pre>
List network hosts and DNS servers.	<pre>CLI@nexenta> net list host CLI@nexenta> net list dns</pre>

Changing Management Address

You must have set up a static IP for the management address during the installation of the appliance but if you need to change the management address after the installation run the following command:

```
CLI@nexenta> config set system.managementAddress=<IP>
```

Verify the changes made to the management IP address by running the following command.

```
CLI@nexenta> config get all system.managementAddress
NAME PROPERTY VALUE
system.managementAddress description Management IP address
system.managementAddress flags --
system.managementAddress name system.managementAddress
system.managementAddress schema <STRING>
```

```
system.managementAddress value 0.0.0.0
```

Managing Users

After installing and rebooting the NexentaStor appliance, you can log in using the admin username and the password you set up during installation.

- ❖ *To change the admin user password:*

```
CLI@nexenta> user passwd -p <new password> admin
```

- ❖ *To create a new user:*

1. You can create as many NexentaStor user accounts as needed but each User you create can be only assigned to a single role, and inherits the access privileges granted to the role. When creating a new user, you can specify the group the user belongs to. If you do not specify the group, the new user you create will belong to a group called Other.

```
CLI@nexenta> user create [options] [-p <password>] [-g <group>] <login>
```

Example:

```
CLI@nexenta> user create -p 1234@abcd testuser
```

2. Verify the user creation.

```
CLI@nexenta> user list
```

LOGIN	UID	PRIMARYGROUP	COMMENT
nobody	60001	nobody	NFS Anonymous Access User
admin	100	other	-
guest	101	other	-
testuser	102	other	-

3. To query user groups:

```
CLI@nexenta> group list
```

NAME	GID	USERS
nobody	60001	nobody
noaccess	60002	noaccess
nogroup	65534	nobody4
other	1000	testuser

4. To create a new user group called allaccess:

```
CLI@nexenta> group create <group name>
```

Example:

```
CLI@nexenta> group create allaccess
```

5. Verify if the group is created.

```
CLI@nexenta> group list
```

```

NAME      GID    USERS
nobody    60001  nobody,testuser1
noaccess  60002  noaccess
nogroup   65534  nobody4
allaccess 100    -

```

6. To add a member to a user group:

```

CLI@nexenta> user create -p 1234@abcd -g allaccess testuser2

NAME      GID    USERS
nobody    60001  nobody,testuser1
noaccess  60002  noaccess
nogroup   65534  nobody4
allaccess 100    testuser2

```

Also see the following commands to manage users and user groups:

```

CLI@nexenta> user -h
CLI@nexenta> group -h

```

Note:

Passwords should be at least 9 characters long and contain at least 3 of the following classes of characters: lowercase, uppercase, numeric, and special (for example, !, @, #,\$, %, ^). Passwords should not be based on English dictionary or slang words, nor English first names or surnames..

Configuring SMTP Email Service

This section demonstrates how to set up SMTP mail server for NexentaStor. Many NexentaStor tasks, such as system failure notification, require that you properly configure the SMTP mail server.

The following table lists the parameters you configure for an SMTP server. Have this information ready before you begin to configure the SMTP server

Table 1-1: SMTP Parameters

SMTP property	
SMTP host	SMTP server hostname or IP address.
SMTP Server Port	Port for the SMTP server. Port 25 is the default, but some IPs deny its use due to the extensive spam and malware traffic this port receives. You can choose to configure another port.
SMTP User Name	Name that you use to access your e-mail. Login name for this SMTP server.
SMTP Password	Password that you use to access your e-mail. Password for the SMTP server login.

SMTP property	
Sender email address	The mail address displayed in the Sender field for all emails originating from NexentaStor.
SMTP Authentication method	Method of SMTP authentication mechanism that your mail server uses. The options are: Plain (default value) SSL TSL
Local superuser admin email	Email address of the superuser admin, required for the recovery of lost passwords.

1. Query and modify the email address of your NexentaStor admin user:

```
CLI@nexenta> config list system.administratorEmail
NAME                               FLAGS  VALUE
system.administratorEmail  --    ''
CLI@nexenta> config set system.administratorEmail=<email address>
```

2. Configure email notifications for alerts by querying and modifying the alert settings:

```
CLI@nexenta> config get all alert
NAME                               FLAGS  VALUE
alert.email.address  --    ''
alert.email.subscribe --    []
```

3. To turn on email alerts, set the email address.

```
CLI@nexenta> config set alert.email.address=abcd@nexenta.com
Now all alerts will be forwarded to the email address specified here.
```

4. You may also limit the alerts only to a specific group of alerts.

```
CLI@nexenta> config set alert.email.subscribe=nef.stat.zpool.usage
This sends only the "pool" usage alters.
CLI@nexenta> config set alert.email.subscribe=defect.sunos,fault.fs.zfs
This sends only the software defect alerts.
```

5. Query and modify SMTP service attributes:

```
CLI@nexenta> svc get smtp
CLI@nexenta> svc set <property>=value smtp
```

Example:

```
svc set host=smtp.gmail.com smtp
svc set port=465 smtp
svc set useSsl=True smtp
svc set useTls=True smtp
svc set authMethods=PLAIN smtp
svc set user=abcd@gmail.com smtp
```

```
svc set password=nexenta1 smtp
```

6. Verify the properties you set using the following command:

```
CLI@nexenta> svc get smtp
Name                Value
authMethods         ['PLAIN']
host                 smtp.gmail.com
password             *****
port                 465
rejectUnauthorized  True
senderEmail         -
useSsl               True
useTls               True
user                 abcd@gmail.com
```

7. To test sending email to your administrator using current SMTP settings:

```
CLI@nexenta> system smtp-test
```

What Comes Next?

In the following chapter you learn how to provision NexentaStor appliances and manage them in your environment.

Configuring NexentaStor 5.0 Appliances

This chapter covers the following topics:

- [Summary of Storage and Cluster Configuration](#)
- [Verifying Enclosure and Disk Information](#)
- [Managing Network](#)
- [Configuring HA Cluster](#)
- [Configuring a Pool](#)
- [Configuring HA Service for the Shared Pool](#)
- [Creating and Sharing a File System](#)
- [Creating and Sharing Volumes](#)
- [What Comes Next?](#)

Summary of Storage and Cluster Configuration

The following table lists the management tasks that you can perform on a NexentaStor appliance(s). Using NexentaStor CLI, you can create a pool on a clustered or a non-clustered appliance, create filesystems and share them for anonymous access or authenticated access in workgroup mode or domain. For a clustered appliance, we can configure pools with a High Availability (HA) service for the pools to failover to the alternate node automatically when the cluster detects a system failure.

To start provisioning a NexentaStor appliance using NexentaStor CLI, complete the tasks in the order presented in this table.

Table 2-1: Task Map

Task	Instructions
<p>Identify disks associated with a NexentaStor appliance</p> <p>Once you have JBODs attached to the appliance, you may want to review the disks that can be assigned as data, cache, spare, or log devices.</p> <p>This section covers the steps to:</p> <ul style="list-style-type: none"> • view all the enclosures and disks enclosed in a chassis. 	<p>See Verifying Enclosure and Disk Information</p>

Task	Instructions
<p>Update NexentaStor appliance's network (if needed)</p> <p>This section covers the steps to:</p> <ul style="list-style-type: none"> • verify that the network interface you configured during NexentaStor appliance installation is in place, • configure aggregates and VLANs to optionally maximize the network throughput and performance; monitor their status. 	<p>See Managing Network</p>
<p>Configure a pool for a single node appliance or clustered appliance</p> <p>Pools are the containers for all datasets such as filesystems, volume groups, and volumes.</p> <p>This section covers the steps to:</p> <ul style="list-style-type: none"> • identify the disks that can be used for creating a pool, • create a pool with the desired redundancy characteristics, • add cache devices, and disk logs to optimize performance, • add spares to improve availability, • add unmap support for SSDs to efficiently use the storage, • schedule a scrub service to check the pool integrity. 	<p>See Configuring HA Cluster and Configuring a Pool</p>
<p>Configure a shared pool HA service for pools to fail over (HA Cluster)</p> <p>The primary benefit of an HA Cluster is for NexentaStor to detect storage system failures and transfer ownership of the shared pools to the alternate NexentaStor appliance.</p> <p>This section covers the steps to:</p> <ul style="list-style-type: none"> • add a shared pool to a HA service, • configure a VIP for the clients to access. 	<p>See Configuring a Pool, and Configuring HA Service for the Shared Pool</p>
<p>Configure a filesystem</p> <p>The filesystem is managed by multiple properties that enable the user to achieve maximum performance and optimization.</p> <p>This section covers the steps to:</p> <ul style="list-style-type: none"> • create a filesystem, • manage a filesystem. 	<p>See Creating and Sharing a File System</p>

Task	Instructions
<p>Share a Filesystem</p> <p>NexentaStor consolidates a number of advanced capabilities to share a filesystem(s) over the network.</p> <p>This section covers the steps to:</p> <ul style="list-style-type: none">• share a filesystem using various sharing protocols (SMB, NFS), manage the Access Control Lists for the shares.	See Creating and Sharing a File System
<p>A volume group is a container for managing volume datasets.</p> <p>These sections demonstrate how to do the following:</p> <ul style="list-style-type: none">• Creating and managing volume groups and volumes• Sharing iSCSI volumes• Sharing to any host on the network• Sharing to a specific set of iSCSI initiators• Preparing iSCSI LUNs for fail over• Sharing FC volumes• Sharing to a specific set of FC initiators• Preparing FC LUNS for fail over	See Creating and Sharing Volumes

Verifying Enclosure and Disk Information

Before creating a new pool, you may want to identify the disks to be assigned as data, cache, or log devices. If you have a JBOD attached to a NexentaStor appliance, you can use the following commands to list all the disks belonging to an appliance to verify their health status and utilization.

1. View the enclosure.

```
CLI@nexenta> enclosure list
CHASSISID  MODEL   LINKRATE LABEL TOTALBAYCOUNT USEDDBAYCOUNT UNITSERIAL
0          LSI SAS616x - - 0 0 0700213V
LEGACY_SATA Nexenta Legacy SATA JBOD - - 6 6
```

2. View the details of a specific enclosure.

```
CLI@nexenta> enclosure get all <CHASSIS ID>
```

Example:

```
CLI@nexenta> enclosure get all LEGACY_SATA
NAME          PROPERTY                VALUE
LEGACY_SATA  chassisId              LEGACY_SATA
LEGACY_SATA  label                  -
LEGACY_SATA  model                  Nexenta Legacy SATA JBOD
LEGACY_SATA  productId              Legacy SATA JBOD
LEGACY_SATA  revision               v.1
LEGACY_SATA  modelInfo.stringData   -
LEGACY_SATA  vendorId               Nexenta
LEGACY_SATA  modelInfo.vendorSpecific -
LEGACY_SATA  sesNodeId              0
LEGACY_SATA  totalBayCount          6
LEGACY_SATA  unitDesignator         -
LEGACY_SATA  unitSerial             -
LEGACY_SATA  usedBayCount           6
```

3. View the sensors for a given enclosure.

```
CLI@nexenta> enclosure sensor <CHASSIS ID>
```

4. View the given disks along with their health status and utilization. When no pool names are given, all disks in the system are listed. The MEDIATYPE column indicates whether the disk is HDD or SSD. An NVME device is listed as a block device (BLKDEV) of type SSD.

```
CLI@nexenta> disk list
```

System response:

NAME	LABELSIZE	MEDIATYPE	STATE	LOCATION	ENCLOSURELABEL	USEDNAME
c2t0d0 -	13.00G	hdd	ONLINE	LEGACY_SAS:0	-	rpool
c2t1d0 -	2.00G	hdd	ONLINE	LEGACY_SAS:1	-	-
c2t2d0 -	2.00G	hdd	ONLINE	LEGACY_SAS:2	-	-

5. View properties for a given disk. These properties can be modified using the disk set subcommand.

```
CLI@nexenta> disk get all <disk name>
```

Example:

```
CLI@nexenta> disk get all c2t0d0
```

Managing Network

During NexentaStor 5.0 installation, you may have set up a network interface card (NIC) for the NexentaStor system. The tasks in this section demonstrate how to verify that the network interface is in place and configured, add a second network interface, use aggregation and VLANs to maximize network throughput and performance.

Verifying the Available Network Interface

During NexentaStor 5.0 installation, you set up the network interface card (NIC) for the NexentaStor appliance. This section demonstrates how to verify the network interface is in place and properly configured. You can view existing network (hardware and software) interfaces (NICs), such as aggregations with the link status.

1. Query the list of network interface cards (NIC) and verify each card's status. Also, note the NIC name that you will need for the next step.

```
CLI@nexenta> link list
NAME      CLASS    STATE    OVER      MTU    SPEED
e1000g0   phys     up       e1000g0   1500   1000
e1000g1   phys     unknown  e1000g1   1500   0
```

2. View the details of the NIC.

```
CLI@nexenta> link get all e1000g0
```

Assigning Addresses to Network Interfaces and Servers

1. Add a physical network interface card if needed. To add an interface, verify that the new NIC is available on the system.

```
CLI@nexenta> inventory nic
NAME      MEDIATYPE STATE    SPEED  MAC                DUPLEX
e1000g0   Ethernet  up       1000   0:25:90:61:b0:14  full
e1000g1   Ethernet  unknown  0      0:25:90:61:b0:15  half
```

2. Configure the newly added interface. The following example assigns a static IP address to the e1000g1 NIC.

```
CLI@nexenta> ip create static <name> <address>
```

Example:

```
CLI@nexenta> ip create static e1000g1/v4 10.3.10.38/24
```

3. Or assign a dynamic IP address to the interface e1000g1 using the following command:

```
CLI@nexenta> ip create dhcp e1000g1/v4
```

4. View the NIC configuration to verify the newly added interface.

```
CLI@nexenta> ip list
```

System response:

```
NAME          TYPE      STATE  ADDRESS
e1000g0/v4   static   ok     10.3.10.38/24
e1000g1/v6   addrconf ok     fe80::20c:29ff:fe8d:82a6/10
```

5. If you have already assigned an IP address and need to update the IP address, destroy the IP address and then assign a new IP as in the previous command.

```
CLI@nexenta> ip destroy e1000g0/v4
```

6. Optionally, set a new DNS name server.

```
CLI@nexenta> net create dns <address>
```

Example:

```
CLI@nexenta> net create dns 10.3.40.245
```

7. Verify if the dns is created successfully.

```
CLI@nexenta> net list dns 10.3.40.245
```

```
NAMESERVER  PROTOCOL
10.3.40.245  ipv4
```

8. Verify the IP configuration, then verify that the NIC state is OK.

```
CLI@nexenta> ip list
```

```
NAME          TYPE      STATE  ADDRESS
e1000g0/v4   static   ok     10.3.10.38/24
e1000g0/v6   addrconf ok     fe80::225:90ff:fe61:b014/10
```

Creating a Network Route

During the NexentaStor 5.0 installation, you must have set up the network route. You can also create a static route manually using the CLI. When creating a route, a network or host should be specified as the destination, and a reachable network address should be specified as the gateway.

1. To create the default route, specify 'default' for the destination.

```
CLI@nexenta> route create <destination> <gateway>
```

Example:

```
CLI@nexenta> route create default 10.3.10.1
```

2. Verify the route you created.

```
CLI@nexenta> route list default
```

```
DESTINATION  GATEWAY    PROTOCOL  INTERFACE  REFS  USE    STATIC
default      10.3.53.1  ipv4      e1000g0    8     228397  yes
```

Depending on your network workload, you may benefit from aggregating network connections or by using virtual LANs (VLANs). If you need to aggregate network connections, ensure that you have two or more NICs available on the system.

Aggregating NICs

Aggregating network links in NexentaStor 5.0 is often used to increase bandwidth within a physical network setup. Rather than being limited to the bandwidth of the largest NIC, you can increase the throughput to that of the combined NICs.

1. To aggregate links in NexentaStor, provide the type of link (aggr), the NIC names separated by a comma, and the aggregate link name (aggr1).

```
CLI@nexenta> link create aggr [-P <policy>] [-L <mode>] [-T <timer>] [-u <mac>] <name> <link>...
```

Example:

```
CLI@nexenta> link create aggr aggr1 e1000g1,e1000g2
```

2. Verify the created aggregation.

```
CLI@nexenta> link list aggr aggr1
NAME      CLASS  STATE OVER MTU   SPEED
aggr1     aggr   up   e1000g1,e1000g2 1500 0
```

Using VLANs

VLANs (Virtual LANs) allow you to group network hosts together according to resource needs, rather than being limited to grouping hosts that are on the same network switch. Using VLANs with NexentaStor 5.0 enables you to have more flexibility in managing and responding to network traffic needs.

1. To set up a VLAN for use with NexentaStor, provide the type of link (vlan), one VLAN link (other links will be ignored), the VLAN ID, and the VLAN name.

```
CLI@nexenta> link assign vlan <name> <vid> <link>
```

Example:

```
CLI@nexenta> link assign vlan vlan1 12 e1000g1
```

2. Verify the VLAN you created.

```
CLI@nexenta> link list vlan vlan1
NAME  OVER  VID
vlan1 e1000g2 12
```

If you are planning to set up a cluster environment for the high availability of the storage pool, configure HA cluster as shown in the section below [Configuring HA Cluster](#). HA cluster runs a defined set of services and monitors each cluster node member for failures and can trigger an automatic failover to the other node in the cluster. The NexentaStor appliances in a cluster are connected through various communication channels and exchange heartbeats that provide information about their states.

Configuring HA Cluster

The term High Availability (HA) Cluster refers to the set of hardware that encompasses a pair of NexentaStor appliances with the shared storage. These clustered nodes are configured with a common VIP and are configured to exchange heartbeats through some communication channels that provide information about their states.

Note: To avail of the HA functionality, the optional HA feature license must be activated on all cluster nodes.

This section uses the following name conventions as examples:

```
ha cluster name      democuster
```

```
1st node name in a  smc-53-109
cluster
```

```
2nd node name in a  smc-53-110
cluster
```

The cluster creation process involves the following tasks:

1. Verify that the HA feature is licensed and enabled on both NexentaStor appliances in the cluster.

```
CLI@nexenta> license show
```

If the HA feature is not activated, contact sales@nexenta.com to obtain your HA license token then run:

```
CLI@nexenta> license activate <key>
```

```
PROPERTY      VALUE
guid          44454c4c-5200-1051-8058-b3c04f563532
valid         yes
status        ok
type          ENTERPRISE
product       NexentaStor
version       5.0.1.0.2
licensee      h.s@acme.com
serial        SR-DEV-NS-201614507
features      fibrechannel, scheduledReplication, highAvailability
```

2. To find out if the HA system feature is running on each cluster node:

```
CLI@nexenta> svc list ha
```

```
NAME  DESCRIPTION          STATE
ha    HA cluster service  online
```

If HA service is not online, enable it:

```
CLI@nexenta> svc enable ha
```

3. Update the hosts to ensure that they resolve to each other. Use the static management IP address.

On the first node in a cluster environment, run the following command:

```
CLI@nexenta> net create host <Static Management IP address of the second node> <host name of the second node>
```

Example:

```
CLI@nexenta> net create host 10.3.53.110 smc-53-110
```

See [Changing Management Address](#)

4. Do the equivalent on the second node.
5. Verify if they resolve to each other by running the following command on both the nodes.

```
CLI@nexenta> net list host
```

Example:

ADDRESS	HOSTNAME	ALIAS	PROTOCOL
::1	localhost	-	ipv6
127.0.0.1	localhost	loghost	ipv4
10.3.53.109	smc-53-109	-	ipv4
10.3.53.110	smc-53-110	-	ipv4

6. Verify that the system day and time settings on each of the cluster member nodes are in sync.
7. Create a cluster by providing the host names of the two nodes in this command. The following command creates a cluster named democluster on the two nodes smc-53-109 and smc-53-110.

```
CLI@nexenta> hacluster create [-fny] [-d <description>] [-H <heartbeats>] <nodes> <cluster>
```

Example:

```
CLI@nexenta> hacluster create -d "test cluster" smc-53-109,smc-53-110 democluster
```

8. Verify the cluster status.

```
CLI@nexenta> hacluster status
```

== Cluster status ==

NAME	STATUS	NODES	SERVICES	DESCRIPTION
democluster	ok	2/2	0/0	No description given

== Nodes ==

NODE	STATUS	SERVICES	ADDRESS	HostId	Release
smc-53-109	up	0/0	10.3.53.109	4ef54012	3.12.0
smc-53-110	up	0/0	10.3.53.110	93c69949	3.12.0

Use the following commands for more details on the CLI cluster commands.

- To list all the HA cluster CLI commands:

```
CLI@nexenta> hacluster
```

- To list the available options for the subcommand:

```
CLI@nexenta> hacluster <subcommand name> -h
```

See the *NexentaStor 5.0 HA QuickStart Guide* in <https://nexenta.com/products/downloads/nexentastor> for more details on how to create, manage, and monitor HA clusters.

Configuring a Pool

Pools are the containers for all datasets such as file systems, volume groups, and volumes. The following sections demonstrate how to create a pool, then add redundancy, cache devices, and disk logs.

This section uses the following name conventions as examples:

```
pool (manual)          poola
pool (auto)            poola-auto
ha service name       poola
1st node name in a   smc-53-109
cluster
2nd node name in a   smc-53-110
cluster
```

Creating a Pool

Before you create a pool, identify the disks to be included in the pool.

Note: If you have a clustered environment and want the pool to be part of the cluster service, you must create a shared pool on one of the HA nodes using shared devices that are visible from both the nodes.

List the disks that you want to include in the pool. Note the logical device names located in the left column of the command output.

```
CLI@nexenta> disk list
NAME      LABELSIZE MEDIATYPE STATE LOCATION ENCLOSURELABEL USEDNAME
c2t0d0 -   13.00G hdd      ONLINE LEGACY_SAS:0 -          rpool
c2t1d0 -   2.00G hdd      ONLINE LEGACY_SAS:1 -          -
c2t2d0 -   2.00G hdd      ONLINE LEGACY_SAS:2 -          -
```

Create the storage pool manually (using the pool create subcommand), or automatically create new pools based on specified options and arguments (using the pool create-auto subcommand). The create-auto subcommand streamlines the process of setting up a NexentaStor appliance, and is the recommended method for creating pools. Assess your site's configuration prior to using this subcommand to ensure the creation of pools that optimize the design.

When creating pools, required options include data redundancy type (for example, RAIDZ1, RAIDZ2, or mirror), the data devices to be included in the pool, and the pool name.

Note: You cannot change the redundancy type after a pool is setup. However, you can add spares and mirrors as shown in the next section.

The following steps can be used for creating a pool on a device shared between the clustered nodes or on a stand alone device that is part of only one node.

1. Before creating a pool, verify if any pool exists on the node.

```
CLI@nexenta> pool list
NAME    SIZE    ALLOC  FREE   AVAIL  DEDUP  EXPANDSZ  FRAG  HEALTH
rpool  12.47G  6.98G  5.50G  44%    1.00x  -          15%  ONLINE
```

2. To create a pool manually on a shared device:

```
CLI@nexenta> pool create [-fnv] [-R altroot] [-o <properties>] <pool>
<vdev>...
```

Example:

```
CLI@nexenta> pool create poola c2t1d1
```

3. Optionally, to enable automatic pool expansion when the associated device is grown.

```
CLI@nexenta> pool set autoexpand=yes <pool name>
```

Example:

```
CLI@nexenta> pool set autoexpand=yes poola
```

Nexenta recommends planning for pool sizes to accommodate future expansion before creating a pool, rather than reactively replacing multiple disks to increase pool capacity as the needs arise.

4. Verify if the autoexpand property is set to “yes”.

```
CLI@nexenta> pool get autoexpand poola
NAME  PROPERTY  VALUE
poola autoexpand yes
```

5. Or, the pool create-auto command will automatically create pool data vdevs based on user provided criteria. For example, you can ask for the automated creation of a mirrored pool pool_m of 120 disks, with 2 way mirrors, across 3 enclosures, with 2TB disks. The system will take care of laying out all the vdevs in the right way and automate the entire setup.

```
CLI@nexenta> pool create-auto <redundancy> <pool> -M <max-devices> [-c
<vdev-size>] [-t <media-type>] [-s <disk-size>] [-r <rpm>] [-N] [-e
<enclosures>] [-R altroot] [-o <properties>] [--config-output=<flags>]
```

Example:

```
CLI@nexenta> pool create-auto mirror poola-auto -c 2 -e e1,e2,e3 -M 120
-t hdd -s 2T
```

Note: When creating a pool, you can force to utilize the disks even if they are currently in use by using the `-f` flag.

6. Verify the pool you created.

```
CLI@nexenta> pool list <pool name>
```

Example:

```
CLI@nexenta> pool list poola
NAME  SIZE  ALLOC  FREE  AVAIL  DEDUP  EXPANDSZ  FRAG  HEALTH
poola 1.92G 100.1M 1.82G 95%    1.00x  -          0%   ONLINE
```

7. To display the list of pool attributes that can be modified:

```
CLI@nexenta> pool set --help
```

If you created a shared pool on one of the nodes in a clustered environment and want to provide high availability to the user's data in case of a node failure, create a HA service as shown in the section [Configuring HA Service for the Shared Pool](#).

Adding Spares, Cache Devices, and Log Disks

You can add spares, cache devices, and log disks using the pool add command.

Adding Spares

To add spares, use the following commands:

1. Identify the disks that are available. Note the logical device names on the left column.

```
CLI@nexenta> disk list
```

2. Add a spare to a pool.

```
CLI@nexenta> pool add <pool name> spare <disk-list>...
```

Adding Cache Devices

To add cache devices, use the following commands:

1. Identify the disks that you have available. Note the logical device names located in the left column.

```
CLI@nexenta> disk list
```

2. Add a cache device named c2t5d0 to the pool named poola.

```
CLI@nexenta> pool add <pool name> cache <disk-list>
```

Example:

```
CLI@nexenta> pool add poola cache c2t5d0
```

Adding Log Disks

To add log disks, use the following commands:

1. Identify the disks that you have available. Note the logical device names located in the left column.

```
CLI@nexenta> disk list
```

2. Add a log device with device name c2t6d0 to the pool named poola.

```
CLI@nexenta> pool add <pool name> log <disk name>
```

Example:

```
CLI@nexenta> pool add poola log c2t6d0
```

Smart-Sparing

The smart-sparing and auto-replace features are available to improve storage availability and simplify maintenance operations. With smart-sparing, the failure of a device (either HDD or SSD) in a pool will automatically activate a spare. Smart-sparing allows a failed HDD to be spared with an SSD but not vice-versa because that would cause performance degradation. Once the spare activation is completed, the user will be able to easily swap out and replace the failed physical device in the pool.

The scope of smart sparing is local or on a single appliance, not global or cluster-wide.

Maintaining Pools (Scrub and Trim/Unmap)

NexentaStor 5.0 supports pool scrubbing and trimming.

Scrub

The scrub process traverses the data of the entire pool and checks to make sure that there are no data integrity issues.

To trigger the start and stop of the Scrub feature for a pool:

```
CLI@nexenta> pool start-scrub
```

```
CLI@nexenta> pool stop-scrub
```

Use a cron expression to set the Scrub schedule. In the example below, a Scrub operation is triggered at 11:15 (PM) every Sunday.

```
CLI@nexenta> pool set scrubSchedule="15 23 * * 7" poola
```

The following graphic illustrates the cron expression structure.

```

┌────────── min (0 - 59)
| ┌────────── hour (0 - 23)
| | ┌────────── day of month (1 - 31)
| | | ┌────────── month (1 - 12)
| | | | ┌────────── day of week (0-7)
| | | | |
| | | | |
| | | | |
15 23 * * 7

```

```
CLI@nexenta> pool get scrubSchedule poola - displays back the Scrub
schedule you configured.
```

Trim

When enabled for pools on SSDs, the Trim/ Unmap feature in NexentaStor notifies the underlying storage media about certain sectors that are no longer needed in a volume or a file system, and thus can be de-allocated made available for other LUNs to use.

You can trigger the start and stop of the Trim feature for a pool using the following commands.

```
CLI@nexenta> pool start-trim
```

```
CLI@nexenta> pool stop-trim
```

The following example demonstrate how to create, stop, and verify a Trim schedule using a cron expression.

```
CLI@nexenta> pool set trimSchedule="0 1 * 2 *" poola
```

```
CLI@nexenta> pool set trimStopSchedule="0 1 * 2 *" poola
```

```
CLI@nexenta> pool get trimSchedule poola
```

Note: Not all SSDs support the Trim option. Verify that your SSD supports this feature.

Configuring HA Service for the Shared Pool

The term HA service refers to the service that runs on the clustered nodes to provide high availability access to the user's data. When the HA service detects a node failure, it transfers ownership of the shared storage to the other node in the cluster pair.

To create a cluster service for the shared pool, you must have created the pool on one of the HA nodes using shared devices that are visible from both the nodes.

Note: HA services can be created only from imported pools. Pools must be explicitly imported on the node that initiated the service creation request.

Nexenta recommends that you always verify that the shared pool is in a healthy state before creating an HA service. The HA cluster command does not detect disk and pool failures.

```
CLI@nexenta> pool status <pool name>
```

The HA service name must match the pool name, a required argument to the command.

You can create an HA service with or without virtual IP addresses (VIPs). Nexenta recommends to create VIP(s) for nfs, smb shares and iSCSI LUNs to successfully failover in case of a disaster.

1. Discover unclustered pools that are available to become members of a cluster.

```
CLI@nexenta> hacluster find-pools
NAME          GUID
poola         1797983880284287424
poola-auto    5795983885284640389
```

2. We can now add a HA service named "poola" to control the currently imported pool (also named poola) on node smc-53-109. When creating a service, you can specify the node where you intend the service to start. You can also specify the list of nodes on which the HA service is allowed to run. If you do not provide the node name, the service is initiated from a random node among available cluster nodes.

The example below also creates a VIP (10.3.53.111/255.255.255.0 with the subnet mask) using the interface e1000g0 from the node smc-53-109 and e1000g0 from the node smc-53-110.

```
CLI@nexenta> haservice create [-evnf] [-d <description>] [-m <node>] [-r <timeout>] [-i <delay>] [-g <guid>] [-N <nodes>] [-V <vips>] <pool>
```

Example:

```
CLI@smc-53-109> haservice create -V vip01@10.3.53.111/255.255.255.0=smc-53-109:e1000g0,smc-53-110:e1000g0 -m smc-53-109 poola
```

3. Optionally, add pools to an HA Service.

When you add a second pool to the HA service, both the pools are now associated with the same poola service and will fail over together as members of the cluster.

```
CLI@nexenta> haservice add-pool <service> <pool> <guid>
```

Example:

```
CLI@nexenta> haservice add-pool poola poola1 <guid>
```

4. Now initiate the HA service.

After you created an HA service, you can set the specific service to start either in auto (-a) or manual (-m) mode after the service stops. The command below specifies that the `poola` service will be started automatically.

```
CLI@nexenta> haservice set-mode [-amnv] <service> <node>
```

Example:

```
CLI@smc-53-109> haservice set-mode -a poola smc-53-109
```

If you selected the manual mode, start the service manually with the following CLI command.

```
CLI@nexenta> haservice start <service> <node>
```

Example:

```
CLI@smc-53-109> haservice start poola smc-53-109
```

5. Verify the service you created.

```
CLI@smc-53-109> haservice status
```

Command output:

```
==service==
poola
==status==
NODE          STATUS      MODE          UNBLOCKED
smc-53-109    running    automatic     yes
smc-53-110    stopped    automatic     yes
==pools==
NAME  GUID          PRIMARY
poola 13660110209781822772 yes
==VIP==
NAME          ADDRESS          IPv6  NODE          NIC
vip01         10.3.53.111/255.255.255.0  no    smc-53-109    e1000g0
smc-53-110    e1000g0
init timeout: 20
run timeout: 8
disk heartbeats: not available
==SCSI reservations==
NODE  DISK  TYPE
universe c2t2d0 SCSI2
```

6. In the event of either node failing (smc-53-109 failing in this example) the surviving node, smc-53-110, takes over the HA service for the pool if it is built using the shared storage devices that are accessible from both the nodes.

You can also trigger a fail over of the HA service(s) running on a node to the other node in the cluster during the maintenance period or if you need to upgrade a node that runs the cluster service. To fail over the HA service use the following command:

```
CLI@nexenta> haservice failover <from-node> <to-node>
```

This command fails over all the services running on the node and imports all the pools to the other node.

7. Verify that the pool from smc-53-109 imported to smc-53-110.

```
CLI@smc-53-110> pool list
```

NAME	SIZE	ALLOC	FREE	AVAIL	DEDUP	EXPANDSZ	FRAG	HEALTH
poola	9.63G	100.2M	9.53G	99%	1.00x	-	0%	ONLINE
rpool	12.47G	7.41G	5.06G	41%	1.00x	-	26%	ONLINE

8. When the failed node is repaired and restarted, it rejoins the cluster. Verify that the recovered node rejoined the cluster.

```
CLI@smc-53-110> hacluster status
```

NODE	STATUS	SERVICES	ADDRESS	HostId	Release
smc-53-109	up	0/1	10.3.53.109	-	-
smc-53-110	up	1/1	10.3.53.110	808b556b	3.12.0

9. Now you as an administrator can control where the pool is redistributed. Now to fail back the pool and the HA service to the repaired node smc-53-109 so that the cluster is back in its original configuration, run the following command on the smc-53-110 node.

```
CLI@nexenta> haservice move <service> <node>
```

Example:

```
CLI@smc-53-110> haservice move poola smc-53-109
```

10. Now verify from both the nodes that the HA service is up and running from their original configuration. Run the following commands from both the nodes.

```
CLI@smc-53-109> haservice status
```

```
CLI@smc-53-110> haservice status
```

```
CLI@smc-53-109> haservice list
```

NAME	VIPs	NODES	RUNNING	STOP
poola	vip01	smc-53-109,smc-53-110	smc-53-109	smc-53-110

For more details on how to manage the HA service:

- ❖ To list all the HA service CLI subcommands:

```
CLI@nexenta> haservice
```

- ❖ To list the available options for the subcommand:

```
CLI@nexenta> haservice <subcommand name> -h
```

See the *NexentaStor 5.0 HA Admin Guide* in <https://nexenta.com/products/downloads/nexentastor> for more details on how to create, manage, and monitor HA services.

Creating and Sharing a File System

Using NexentaStor 5.0, you can create and share file systems within the pools. To begin, you first create a pool that is always the root directory in the ZFS file system hierarchy. You can set file system properties such as compression mode and reservation size.

This section uses the following name conventions as examples

pool	poola
filesystem for nfs VMware share	poola/testnfs-vmware
filesystem for smb share	poola/testnfs
filesystem for open share	poola/testnfs-open

Creating a File System

This section demonstrates the steps involved when creating a file system.

1. Create the file system that you will share.

```
CLI@nexenta> filesystem create <filesystem>
```

Example:

```
CLI@nexenta> filesystem create poola/testnfs
```

2. See what other file system attributes you can set such as compressionMode (default value is lz4).

```
CLI@nexenta> filesystem set --help
```

With smart compression, compression efficiency is monitored and automatically stops compressing blocks if there are no significant space savings.

3. Verify that the file systems have been created.

```
CLI@nexenta> filesystem list poola/testnfs
```

PATH	USED	AVAIL	REFER	NFS	SMB	MOUNTPOINT
poola/testnfs	1.55T	5.03T	593.19G	no	no	/poola/testnfs

4. View all the properties of the file system you created.

```
CLI@nexenta> filesystem get all poola/testnfs
```

5. To see currently mounted file systems and their mount points:

```
CLI@nexenta> filesystem mount
```

PATH	MOUNTPOINT
poola/testnfs	/poola/testnfs

Sharing a File System via SMB

SMB service is disabled by default. However, when you create an SMB share, the service gets enabled automatically. NexentaStor 5.0 supports SMB 1.0, SMB 2.1(default version), and SMB 3.0.

Create an SMB Share in Workgroup Mode with Guest Access

By default, a guest account is configured as part of the NexentaStor 5 installation, but SMB access is disabled. To share a file system via SMB in workgroup mode with Guest access enabled:

1. Confirm that guest account exists:

```
CLI@nexenta> user list
LOGIN      UID      PRIMARYGROUP  COMMENT
nobody     60001   nobody        NFS Anonymous Access User
noaccess   60002   noaccess      No Access User
nobody4    65534   nogroup       SunOS 4.x NFS Anonymous Access User
admin      100     other         -
guest      101     other         -
```

2. Enable Guest access for the SMB service:

```
CLI@nexenta> svc set enableGuest=true smb
CLI@nexenta> svc get smb
Name                                     Value
systemComment                           -
restrictAnonymous                       True
preferredDomainController                -
signing                                  enabled
enableNetbios                            False
lanManagerAuthLevel                     4
traverseMounts                          True
keepAlive                                5400
activeDirectorySite                      -
maxProtocolVersion                       2.1
enableIpv6                                True
maxWorkers                               1024
enableDdns                               False
enableGuest                              True
```

3. Create file system poola/test smb:

```

CLI@nexenta> fs create poola/test smb
CLI@nexenta> fs list
PATH            USED      AVAIL      REFER  NFS  SMB  MOUNTPOINT
poola           102.1M   153.94G   26.6K  no   no   /poola
poola/fs1       1.7M    153.94G   1.7M   yes  no   /poola/fs1
poola/fs2       25.3K   153.94G   25.3K  no   no   /poola/fs2
poola/test smb  25.3K   153.94G   25.3K  no   no   /poola/test smb

```

4. Share the file system via SMB and enable Guest access on that share. The following example shows how to set a custom share name “nexenta_share” for the share you create.

```

CLI@nexenta> smb share -o name=nexenta_share poola/test smb
CLI@nexenta> smb set guestok=yes poola/test smb
CLI@nexenta> smb get all poola/test smb
PATH            PROPERTY      VALUE
poola/test smb  access        no
poola/test smb  cache         manual
poola/test smb  guestOk       yes
poola/test smb  description   -
poola/test smb  name          poola_test smb
poola/test smb  quota         yes
poola/test smb  shareState    online

```

5. Note the name of the SMB share that must be used to access it from SMB clients. If you did not provide a custom name for the share, you can always modify the default share name using the following command.

```

CLI@nexenta> smb set shareName=nexenta_share poola/test smb

```

6. Finally, configure the ACL on the file system and add an Access Control Entry (ACE) for the Guest account. By default, a new file system is created with 2 ACEs on it:

```

CLI@nexenta> acl list poola/test smb

```

System response:

INDEX	TYPE	PRINCIPAL	PERMISSIONS	FLAGS
0	allow	owner@	rw xpdDaARWcCos	fd-----
1	allow	groupsid:Administrators@BUILTIN	rw xpdDaARWcCos	fd-----

7. The second ACE allows Administrator accounts from a Windows host to configure user access to that share. To fully enable Guest access to that share, you can also simply add the following ACE:

```
CLI@nexenta> acl set A+groupsid:Guests@BUILTIN:rxpdaARWcs:fd:allow
poola/test smb
```

or

```
CLI@nexenta> acl set A+groupsid:Guests@BUILTIN:modify_set:fd:allow
poola/test smb
```

```
CLI@nexenta> acl list poola/test smb
```

System response:

INDEX	TYPE	PRINCIPAL	PERMISSIONS	FLAGS
0	allow	owner@	rxpdaARWcCos	fd-----
1	allow	groupsid:Administrators@BUILTIN	rxpdaARWcCos	fd-----
2	allow	groupsid:Guests@BUILTIN	rxpdaARWcs	fd-----

The share is now enabled and accessible by Guest accounts at poola_test smb.

Create an SMB Share in Domain Mode

Using Active Directory

To use the DNS-based Active Directory (AD) services, follow these steps:

Note: In many AD environments your AD servers are also DNS servers.

1. Identify the available AD name server, noting its IP address:

```
CLI@nexenta> net list dns
```

2. If no AD name server exists, add two or more AD name servers.

Use this command to add the Active Directory name server if it is not available.

```
CLI@nexenta> net create dns 10.3.40.245
```

3. Verify the domain.

```
NAMESERVER    PROTOCOL
10.3.40.245   ipv4
8.8.8.8       ipv4
```

4. Join the AD domain by specifying your AD administrator username and password, and the AD server name.

```
CLI@nexenta> smb join domain domainadministrator company.corp
```

5. Provide the password when the system prompts for it.

6. Verify the join.

```
CLI@nexenta> smb status
REALMNAME      MODE      DOMAINCONTROLLER
company.corp  domain  company.corp
```

7. Configure and verify NTP before joining the domain.

- Verify that the NTP service is running then synchronize the NexentaStor server time with the domain controller time.

```
CLI@nexenta> svc list ntp
NAME DESCRIPTION STATE
ntp NTP client online
CLI@nexenta> svc set servers=yourntpserver.example.com ntp
```

- Verify that your NexentaStor system has joined the AD domain.

```
CLI@nexenta> smb status
```

Share Filesystem

- Create file system poola/test smb:

```
CLI@nexenta> fs create poola/test smb
CLI@nexenta> fs list
PATH USED AVAIL REFER NFS SMB MOUNTPOINT
poola 102.1M 153.94G 26.6K no no /poola
poola/fs1 1.7M 153.94G 1.7M yes no /poola/fs1
poola/fs2 25.3K 153.94G 25.3K no no /poola/fs2
poola/test smb 25.3K 153.94G 25.3K no no /poola/test smb
```

- Share the file system via SMB. The following example shows how to set a custom share name "nexenta_share" for the share you create.

```
CLI@nexenta> smb share -o name=nexenta_share poola/test smb
```

Sharing a File System via NFS

NexentaStor 5.0 supports NFS versions 3 (default) and 4.

Create an Open NFS Share for a VMware NFS Datastore

To configure an NFS share to be used as a VMware NFS Datastore, we create an NFSv3 share that is open with root access to any ESXi host on the network.

Note: Use this configuration only on a closed network where all hosts are trusted.

- Create file system poola/test nfs-vmware

```
CLI@nexenta> filesystem create poola/test nfs-vmware
CLI@nexenta> fs list
PATH USED AVAIL REFER NFS SMB MOUNTPOINT
poola 102.0M 153.94G 25.3K no no /poola
poola/fs1 1.7M 153.94G 1.7M yes no /poola/fs1
```

```
poola/fs2          25.3K   153.94G  25.3K  no   no   /poola/fs2
poola/testnfs-vmware 25.3K   153.94G  25.3K  no   no   /poola/testnfs-vmware
```

2. To ensure that VMware hosts will negotiate an NFSv3 share and avoid any interoperability problems with NFSv4 negotiations, we set the NFS service protocol maxVersion to 3:

```
CLI@nexenta> svc set maxversion=3 nfs
```

```
CLI@nexenta> svc get nfs
```

```
Name           Value
delegation     True
servers        256
maxVersion     3
gracePeriod    90
nfsMapIdDomain -
minVersion     2
lockdServers   256
```

3. Enable NFS sharing on the file system with root access for any hosts that are trusted on the network:

```
CLI@nexenta> nfs share -o anon=root,securityContexts=sec=sys,rw=* poola/
testnfs-vmware
```

```
CLI@nexenta> nfs get all poola/testnfs-vmware
```

```
PATH           PROPERTY      VALUE
poola/testnfs-vmware anon          root
poola/testnfs-vmware nohide        no
poola/testnfs-vmware securityContexts sec=sys,rw=*
poola/testnfs-vmware shareState     online
poola/testnfs-vmware gidMap           -
poola/testnfs-vmware recursive        -
poola/testnfs-vmware rootMapping      -
poola/testnfs-vmware uidMap           -
```

The NFS share is ready to mount on VMware ESXi hosts.

A more secure option for sharing an NFS share to ESXi hosts is to limit root access to an explicit list of these hosts. The example below disables anonymous access and limits root access to the set of specified IP addresses:

```
CLI@nexenta> nfs share -o anon=-
1,securityContexts=sec=sys,rw=192.168.56.40,192.168.56.41,root=192.168.
56.40,192.168.56.41 poola/testnfs-vmware
```

```
CLI@nexenta> nfs get all poola/testnfs-vmware
```

PATH	PROPERTY	VALUE
poola/testnfs-vmware	anon	-1
poola/testnfs-vmware	nohide	no
poola/testnfs-vmware	securityContexts	sec=sys,root=192.168.56.40,192.168.56.41,rw=192.168.56.40,192.168.56.41
poola/testnfs-vmware	shareState	online
poola/testnfs-vmware	gidMap	-
poola/testnfs-vmware	recursive	-
poola/testnfs-vmware	rootMapping	-
poola/testnfs-vmware	uidMap	-

Create an Open NFS Share for Generic Clients

To configure an NFS share to be used as an open share for generic clients, we will share with anonymous access disabled and add an ACE with appropriate modify permissions for 'everyone' on the file system.

1. Create file system poola/testnfs-open

```
CLI@nexenta> filesystem create poola/testnfs-open
```

```
CLI@nexenta> fs list
```

PATH	USED	AVAIL	REFER	NFS	SMB	MOUNTPPOINT
poola	102.1M	153.94G	26.6K	no	no	/poola
poola/fs1	1.7M	153.94G	1.7M	yes	no	/poola/fs1
poola/fs2	25.3K	153.94G	25.3K	no	no	/poola/fs2
poola/testnfs-open	25.3K	153.94G	25.3K	no	no	/poola/testnfs-open
poola/testnfs-vmware	25.3K	153.94G	25.3K	yes	no	/poola/testnfs-vmware

2. Enable NFS sharing on the file system with anonymous access disabled:

```
CLI@nexenta> nfs share -o anon=-1 poola/testnfs-open
```

If you do not want to set up a security type for your nfs share, use "none" as the security contexts as shown in the following command.

```
CLI@nexenta> nfs share -o securityContexts=sec=sys poola/testnfs-open
```

```
CLI@nexenta> nfs get all poola/testnfs-open
```

PATH	PROPERTY	VALUE
poola/testnfs-open	anon	-1
poola/testnfs-open	nohide	no


```

poola/testnfs-open securityContexts sec=sys
poola/testnfs-open shareState      online
poola/testnfs-open gidMap          -
poola/testnfs-open recursive      -
poola/testnfs-open rootMapping     -
poola/testnfs-open uidMap          -

```

3. Before the NFS share can be mounted by generic clients, we need to add an ACE on the file system giving modify permissions to 'everyone':

```
CLI@nexenta> acl list poola/testnfs-open
```

INDEX	TYPE	PRINCIPAL	PERMISSIONS	FLAGS
0	allow	owner@	rwxpdaARwCcos	fd-----
1	allow	groupsid:Administrators@BUILTIN	rwxpdaARwCcos	fd-----

```
CLI@nexenta> acl set A+everyone@:modify_set:fd:allow poola/testnfs-open
```

```
CLI@nexenta> acl list poola/testnfs-open
```

INDEX	TYPE	PRINCIPAL	PERMISSIONS	FLAGS
0	allow	owner@	rwxpdaARwCcos	fd-----
1	allow	groupsid:Administrators@BUILTIN	rwxpdaARwCcos	fd-----
2	allow	everyone@	rwxpdaARwC--s	fd-----

The file system is ready to mount on any NFS client on the network.

Enabling anonymous access on NFS share:

1. To enable anonymous access on this NFS share, we can set the anon property of the NFS share to 'nobody' and add the appropriate ACE for user:nobody (read set in the example below):

```
CLI@nexenta> user list
```

LOGIN	UID	PRIMARYGROUP	COMMENT
nobody	60001	nobody	NFS Anonymous Access User
noaccess	60002	noaccess	No Access User
nobody4	65534	nogroup	SunOS 4.x NFS Anonymous Access User
admin	100	other	-
guest	101	other	-

```
CLI@nexenta> nfs set anon=nobody poola/testnfs-open
```

```
CLI@nexenta> nfs get all poola/testnfs-open
```

PATH	PROPERTY	VALUE
poola/testnfs-open	anon	nobody
poola/testnfs-open	nohide	no

```

poola/testnfs-open securityContexts sec=sys
poola/testnfs-open shareState      online
poola/testnfs-open gidMap           -
poola/testnfs-open recursive       -
poola/testnfs-open rootMapping      -
poola/testnfs-open uidMap           -

```

```
CLI@nexenta> acl set A+user:nobody:read_set:fd:allow poola/testnfs-open
```

```
CLI@nexenta> acl list poola/testnfs-open
```

INDEX	TYPE	PRINCIPAL	PERMISSIONS	FLAGS
0	allow	owner@	rxpdaARWcCos	fd-----
1	allow	groupsid:Administrators@BUILTIN	rxpdaARWcCos	fd-----
2	allow	everyone@	rxpdaARWc--s	fd-----
3	allow	user:nobody	r-----a-R-c---	fd-----

Creating and Sharing Volumes

In addition to file services, NexentaStor supports both iSCSI and Fibre Channel block services. This section provides an example of commands and steps required to create a volume, create iSCSI targets, target group and host group and share that volume as an iSCSI LUN.

Before we can create a volume, NexentaStor 5 requires the creation of a higher level dataset called a Volume Group. The default dataset hierarchy in NexentaStor 5 is pool/volume-groups/volumes. Volume Groups serve 2 main purposes:

- A VG provides a parent dataset with a customizable set of properties (such as block size or compression settings) that can be simply inherited by all volumes subsequently created in the Volume Group.
- A VG can act as a consistency group for local snapshots or replication snapshots of all the volumes it contains. This can be particularly useful when multiple volumes are configured to support a specific application. Grouping these volumes in the same VG makes taking application consistent snapshots of all the underlying volumes very easy.

Note: Support for VMware vStorage API for Array Integration (VAAI) block is enabled by default.

Creating Volume Groups and Volumes

This section uses the following name conventions:

Volume group	sqlvg
Volume	sql-1
iSCSI Target Group	tg1
host group	hg1

The steps involved in creating volume groups and volumes are:

1. Create a volume group “sqlvg” before you create a storage volume.

```
CLI@nexenta> vg create poola/sqlvg
CLI@nexenta> vg list
PATH          USED   AVAIL   REFER  BLKSIZE
poola/sqlvg  24.0K  153.94G  24.0K  32.0K
```

2. Create a volume in that volume group

```
CLI@nexenta> vol create poola/sqlvg/sql-1 10G
CLI@nexenta> vol list
PATH          VOLSIZE  USED    AVAIL   REFER  BLKSIZE
poola/sqlvg/sql-1  10.00G  10.08G  153.94G  10.7K  32.0K
```

3. Verify that iSCSI Target service is enabled on the appliance

```
CLI@nexenta> svc list
NAME      DESCRIPTION          STATE
```

```

ha      HA cluster service    online
idmap   idmap service            online
iscsit  iSCSI target service    online
ndmp    NDMP backup service    disabled
nfs     NFS server                online
ntp     NTP client                disabled
smb     SMB server              disabled
snmp    SNMP service            online
stmf    STMF service           online
vscan   Virus scan service      disabled

```

Sharing iSCSI Volumes

This section lists the commands related to sharing iSCSI volumes.

1. The first step is to verify that the iSCSI target service is enabled.

```
CLI@nexenta> svc list iscsit
```

2. If the iSCSI target service is not online, enable it.

```
CLI@nexenta> svc enable iscsit
```

3. To query and set iSCSI target service attributes:

```
CLI@nexenta> svc get iscsit
```

```
CLI@nexenta> svc set <property> iscsit
```

4. In order to share our poola/sqlvg/sql-1 volume, we need to create an iSCSI Target Group that will contain one or more iSCSI Targets. In our example, we have 2 IP interfaces that we will configure as portals for 2 different iSCSI Targets:

```
CLI@nexenta> ip list
```

```

NAME          TYPE    STATE  ADDRESS
e1000g0/v4    static ok     192.168.56.5/24
e1000g1/v4    static ok     192.168.56.6/24
e1000g2/v4    dhcp   ok     192.168.0.34/24
lo0/v4        static ok     127.0.0.1/8
lo0/v6        static ok     ::1/128

```

5. Create iSCSI target:

```
CLI@nexenta> iscsitarget create 192.168.56.5
```

New target 'iqn.2010-08.org.illumos:02:e2faf5b9-7b93-403e-f3df-84b5198a19be' has been created.

```
CLI@nexenta> iscsitarget create 192.168.56.6
```

New target 'iqn.2010-08.org.illumos:02:08ba0ff1-2e4f-4b5a-f5ee-f36c765a9644' has been created.

Note: If you have a clustered environment and want the volume to be part of the cluster service, you must create the target using the VIP address you created when configuring the HA service. See [Step 2](#) under [Configuring HA Service for the Shared Pool](#).

6. We then create a single iSCSI Target Group (tg1) containing our 2 iSCSI Targets:

```
CLI@nexenta> iscsitarget list
```

NAME	PORTALS	AUTH	STATE
02:08ba0ff1-2e4f-4b5a-f5ee-f36c765a9644	192.168.56.6:3260	default	online
02:e2faf5b9-7b93-403e-f3df-84b5198a19be	192.168.56.5:3260	default	online

```
CLI@nexenta> targetgroup create tg1 02:08ba0ff1-2e4f-4b5a-f5ee-f36c765a9644 02:e2faf5b9-7b93-403e-f3df-84b5198a19be
```

```
CLI@nexenta> targetgroup list
```

NAME	MEMBERS
tg1	02:08ba0ff1-2e4f-4b5a-f5ee-f36c765a9644, 02:e2faf5b9-7b93-403e-f3df-84b5198a19be

At this point, we can use this iSCSI Target Group to share our volume, either to any host that has access to the iSCSI targets, or to a specific Host Group with an explicit set of iSCSI initiators we want to map our iSCSI LUN to.

Sharing to Any Host on the Network

7. To share to any host on the network:

```
CLI@nexenta> vol list
```

PATH	VOLSIZE	USED	AVAIL	REFER	BLKSIZE
poola/sqlvg/sql-1	10.00G	10.08G	153.94G	10.7K	32.0K

```
CLI@nexenta> targetgroup list
```

NAME	MEMBERS
tg1	02:08ba0ff1-2e4f-4b5a-f5ee-f36c765a9644, 02:e2faf5b9-7b93-403e-f3df-84b5198a19be

```
CLI@nexenta> lunmapping create poola/sqlvg/sql-1 tg1 all
```

```
CLI@nexenta> lunmapping list
```

ID	VOLUME	TARGETGROUP	HOSTGROUP	LUN
0695AFC2686540CF00000000	poola/sqlvg/sql-1	tg1	all	0

Sharing to a Specific Set of iSCSI Initiators

To limit sharing to a specific set of iSCSI initiators, we first create a Host Group hg1 and then the appropriate lunmapping:

```

CLI@nexenta> hostgroup create hg1 iqn.2016-
06.com.mydomain.iscsi:sqlserver
CLI@nexenta> hostgroup list
NAME    MEMBERS
hg1     sqlserver
CLI@nexenta> lunmapping create poola/sqlvg/sql-1 tg1 hg1
CLI@nexenta> lunmapping list
ID                VOLUME                TARGETGROUP  HOSTGROUP  LUN
0695AFC2686540CF00000000 poola/sqlvg/sql-1  tg1          hg1        0

```

Preparing iSCSI LUNs for Failover

For iSCSI shares to successfully failover, ensure that you have the following.

- HA service must be created with VIP(s) as shown in [Step 2](#) under the section [Configuring HA Service for the Shared Pool](#).

Unless you create the VIP(s), the network clients will not be able to access the data in case of a failover.

- If you created a HA service without a VIP and want to add one, use the following command:

```

CLI@nexenta> haservice add-vip poola vip05 9.8.7.6/255.0.0.0 smc-53-
109:e1000g0,smc-53-110:e1000g0

```

In the above example, the VIP vip05 with IP address 9.8.7.6 and netmask 255.0.0.0 is added to the HA service named poola.

- Targets must be configured as explained in [Sharing iSCSI Volumes](#).
- Target groups must be configured as explained in [Sharing iSCSI Volumes](#).
- LUNs are mapped as shown in [Sharing to Any Host on the Network](#) and [Sharing to a Specific Set of iSCSI Initiators](#).
- Verify that you can mount volume shares for clients.
- Check for network connectivity between primary and failover node.

Sharing FC Volumes

When you connect FC ports to a NexentaStor appliance, they are automatically discovered by the system.

Fibre Channel ports operate in the following modes:

- **Initiator** - Enables NexentaStor to access remote FC targets. For example, you can connect LUNs from other NexentaStor appliances. By default, the FC ports are in initiator mode. To change the default mode to target, use the command below:

```
CLI@nexenta> config set system.fcDefaultPortMode=target
```

Reboot the appliance to reflect the changes made to the mode.

- **Target** - Provides the access to NexentaStor FC targets for remote initiators.

The steps required to share a volume as a Fibre Channel (FC) LUN are identical to the steps described above for iSCSI LUN. The only real difference is that Targets and Initiators are identified via World Wide Names (WWNs) instead of IQNs.

1. View the information on the FC HBA to see all available target ports.

```
CLI@nexenta> fctarget list
```

NAME	NODEWWN	CURRSPEED	STATE
wwn.10000000c9bb4b5a	20000000c9bb4b5a	4Gb	online
wwn.10000000c9bb4b5b	20000000c9bb4b5b	4Gb	online

2. Create and name the target group to manage LUN mappings.

```
CLI@nexenta> targetgroup create <name> <FC target in terms of WWNs>
```

The following example shows a targetgroup (tgfc) created with the above WWNs.

Example:

```
CLI@nexenta> targetgroup create tgfc wwn.10000000c9bb4b5a,
wwn.10000000c9bb4b5b
```

3. To verify the target group that was created:

```
CLI@nexenta> targetgroup list
```

NAME	MEMBERS
targetgrp	02:9a4cf019-e94f-6693-ad06-b104be631f7e
tgfc	wwn.10000000c9bb4b5a, wwn.10000000c9bb4b5b

4. To add or remove FC targets in the target group:

```
CLI@nexenta> targetgroup add <targetgroup> <FC target>
```

```
CLI@nexenta> targetgroup remove <targetgroup> <FC target>
```

Sharing to a Specific Set of FC Initiators

1. Create LUN mappings and share the FC LUN to a specific initiator.

To share the FC LUN to a specific initiator, create a host group using one or more initiators.

```
CLI@nexenta> hostgroup create <hostgroup> <wwn of FC initiator>
```

Example:

```
CLI@nexenta> hostgroup create hgfc wwn.21000024FF4899E9
```

2. Verify the created hostgroup.

```
CLI@nexenta> hostgroup list
```

```
NAME MEMBERS
```

```
hgfc wwn.21000024FF4899E9
```

3. Create a volume group and volume as shown in [Creating Volume Groups and Volumes](#) for the FC LUN.

Note: the volume that you share using FC must not be part of other mapping.

4. Map the LUN

```
CLI@nexenta> lunmapping create <volume path> <target group name> <Host group of FC initiators>
```

Where:

<target group> - group of FC targets

<volume-path> - path of the volume to be exposed by this target group

```
CLI@nexenta> lunmapping create poola/sqlvg/sql-1 tgfc hgfc
```

```
CLI@nexenta> lunmapping list
```

ID	VOLUME	TARGETGROUP	HOSTGROUP	LUN
0695AFC2686540CF00000000	poola/sqlvg/sql-1	tgfc	hgfc	0

5. Verify your LUN mappings and LUN list after your mapping is completed.

```
CLI@nexenta> lunmapping list
```

ID	VOLUME	TARGETGROUP	HOSTGROUP	LUN
FD2A4FC33621F5F400000000	fcPool/fcVG/fcV	tgfc	hgfc	0

```
CLI@nexenta> logicalunit list
```

VOLUME	VOLSIZE	WPROTECT	STATE
poola/sqlvg/sql-1	2.79T	no	online

Preparing FC LUNs for Failover

For FC shares to successfully failover, ensure that you have done the following.

1. Enable ALUA on any one of the HA cluster node.

Ensure that you have cluster setup before enabling ALUA.

```
CLI@nexenta> config set ha.alua.enabled=true
```

2. Verify if ALUA is enabled.

```
CLI@nexenta> config list ha.alua
```



```
NAME      FLAGS  VALUE
ha.alua  --      enabled: true
```

3. Targets must be configured as explained in [Sharing FC Volumes](#)
4. Target groups must be configured as explained in [Sharing FC Volumes](#)
5. LUNs are mapped as shown in [Sharing to a Specific Set of FC Initiators](#)

What Comes Next?

In the following chapter you will learn how to protect any dataset, be it a file system, volume group, or volume using some of the features available in NexentaStor 5.0.

Data Protection

This chapter covers the following topics:

- [Protecting Data](#)
- [Configuring Snapshots and Snapshot Schedules](#)
- [Configuring High Performance Replication](#)
- [Managing the Created Replication Service](#)
- [Configuring NDMP for Backups](#)
- [Using Virus Scan Engines](#)
- [What Comes Next?](#)

Protecting Data

Nexenta offers different types of protection services to protect the datasets such as pools, filesystems, volume groups, and volume. This chapter will introduce you to some of the methods available in the appliance to protect the datasets.

List of methods:

- Local Scheduled Snapshot service
A snapshot is a read-only point-in-time representation of a file system, volume that can later be cloned. You can clone a snapshot to create an editable copy.
- Long Distance Replication Service
Long distance replication service formally known as high performance replication service is a dataset protection service that generates snapshots at the primary appliance following a set schedule (Scheduled Replication) or on a continuous basis (Continuous Replication).
- NDMP
Network Data Management Protocol (NDMP) is a networking protocol and an open standard for backing up data in a heterogeneous environment.

Configuring Snapshots and Snapshot Schedules

A snapshot refers to the state of a dataset at a certain point in time and is comprised of a set of reference markers or pointers to data stored on a disk or SAN. You can create and destroy a snapshot, but you cannot modify it. To create an editable copy, clone the snapshot. Snapshot allows you to safely rollback to the previous state after 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.

Creating and Deleting Local Snapshots

To create, list, and delete unreplicated snapshots, see the commands in Table below. Snapshot names are prepended with the filesystem or volume path followed by the @ sign. For example, ABC/fs1@hpr-2016-07-25-07-29-54-041. Table lists basic operations that you can do on a snapshot.

Table 3-1: Creating Snapshots

Task	Related CLI Command
Create a snapshot	<pre>CLI@nexenta> snapshot create [-o <properties>] <snapshot></pre> <p>Example:</p> <pre>CLI@nexenta> snapshot create poola/test smb@snapa</pre> <p>where poola/test smb is the filesystem</p>
Verify properties of a specific snapshot	<pre>CLI@nexenta> snapshot get (all <properties>) <snapshot></pre> <p>Example:</p> <pre>CLI@nexenta> snapshot get all poola/test smb@snapa</pre>
List all snapshots	<pre>CLI@nexenta> snapshot list</pre>
Delete a snapshot	<pre>CLI@nexenta> snapshot destroy <snapshot>...</pre> <p>If a hold exists on a snapshot such that it cannot be destroyed, attempts to destroy the snapshot will return an ERROR. Release the hold on the snapshot using the <code>snapshot release</code> command before destroying the snapshot.</p>
Clone a snapshot	<pre>CLI@nexenta> snapshot clone <snapshot> <clone></pre>
Release a snapshot to be able to destroy it if needed	<pre>CLI@nexenta> snapshot release <tag> <snapshot></pre> <p><tag> represents the name of the hold tag to be destroyed.</p>
View all the available snapshot commands	<pre>CLI@nexenta> snapshot -h</pre>

Creating a Local Scheduled Snapshot Service

You can automatically create a recurring snapshot (snapping job) of a file system or a volume group/volume.

Table 3-2: Scheduling Snapshots

Task	Related CLI Command
Create a recurring snapshot by creating a snapping job	<pre>CLI@nexenta> snapping create [-nrtv] [--description=<desc>] cron=<period> --keep=<n>] <dataset> [<name>]</pre> <p>In the example right below, a snapping job named <code>snaptsr</code> is created for the file system <code>pooltsr/fstsr</code>. Snapshots are to be taken every hour and the three most recent snapshots will be stored.</p> <pre>CLI@nexenta> snapping create --cron='hourly' --keep=3 poola/test smb snaptsr</pre>
Enable the snapping job	<p>To run the snapping job you created with the configured <code>cron</code> schedule, you must enable it first.</p> <pre>CLI@nexenta> snapping enable <snapping job name></pre>
Get list of snapping jobs and their status	<pre>CLI@nexenta> snapping list</pre>
Get the list of snapshots associated with a snapping job	<pre>CLI@nexenta> snapping snapshots <snapping job name></pre>
Run a snapping job on-demand	<pre>CLI@nexenta> snapping run <snapping job name></pre>
Modify a snapping Job	<p>To modify the properties of a snapping job, you must disable it first.</p> <pre>CLI@nexenta> snapping disable <snapping job name> CLI@nexenta> snapping set description="MyCo" <snapping job name></pre>
Add a schedule to a snapping job	<p>A snapping job can have multiple schedules assigned to it. To add a schedule to the snapping job you created, use the following command.</p> <pre>CLI@nexenta> snapping schedule-add [-nv] <snapping job name> <cron> <keep> [<schedule name>]</pre>
Delete a snapping job	<pre>CLI@nexenta> snapping destroy <snapping job name></pre>
Enable or disable a snapping schedule	<pre>CLI@nexenta> hpr schedule-enable <service name> <schedule name> CLI@nexenta> hpr schedule-disable <service name> <schedule name></pre>
List all available snapping commands	<pre>CLI@nexenta> snapping -h</pre>

Managing Snapshots

The following table lists the operations you can perform on a snapshot once it is created.

Table 3-3: Managing Snapshots

Task	Related CLI Command
Clone a snapshot	<p>By definition, snapshots are read-only. You can clone a snapshot to create an editable or writeable copy and then manage this clone separately.</p> <pre>CLI@nexenta> snapshot clone <snapshot> <clone></pre> <p>Where <clone> specifies the new clone file system or volume path.</p>
Rolling back a snapshot	<p>You can rollback a file system or volume group/volume to a snapshot of its previous state.</p> <pre>CLI@nexenta> snapshot rollback [-rdDfNV] <snapshot></pre>
Set properties of a snapshot	<pre>CLI@nexenta> snapshot set [-rNV] <properties> <snapshot></pre> <pre>CLI@nexenta> snapshot reset [-rNV] <properties> <snapshot></pre> <p>Resets a modifiable property for a given snapshot to either the inherited value from an ancestor object, a default value, or the deleted value depending on the semantics of the property.</p>
Retrieve the values of a snapshot	<pre>CLI@nexenta> snapshot get all <snapshot></pre>
Put a hold on a snapshot	<p>Putting a hold on a dataset snapshot prevents it from being destroyed.</p> <pre>CLI@nexenta> snapshot hold [-rNV] < hold tag name> <snapshot></pre> <pre>CLI@nexenta> snapshot holds</pre> <p>lists all the snapshot holds you have configured</p>
Removing a hold	<p>To remove the hold on the snapshot, use the following command:</p> <pre>CLI@nexenta> snapshot release <tag> <snapshot></pre>
Deleting a snapshot	<p>If a hold exists on a snapshot, attempts to destroy that snapshot command will return an ERROR. Release the hold on the snapshot using the <code>snapshot release</code> command before destroying the snapshot.</p> <pre>CLI@nexenta> snapshot destroy [-rNV] <snapshot></pre>
Promoting a cloned snapshot	<p>When a snapshot is cloned, a dependency is created between the clone and the snapshot such that you cannot delete the snapshot for as long as the clone exists. To make a file system or a volume independent of its origin snapshot, you can promote the dataset.</p> <pre>CLI@nexenta> snapshot promote <dataset></pre> <p>Use this command only if the datasets do not belong to a HPR replication service.</p>

Note: You can clone the snapshots that belong to a HPR protection service but do not promote these clones.

Configuring High Performance Replication

About HPR

High Performance Replication (HPR) is a dataset protection service that generates snapshots at the source dataset following a set schedule or on a continuous basis. These snapshots can be stored locally on the source appliance (same appliance where the HPR service was created) or locally and remotely in one or more secondary appliances. Replication is a means to do backup and archive that can be useful during disaster recovery, network failures, or sudden power outages. In the event of a disaster and when the source appliance becomes unavailable, the network clients can access their data from the destination site/appliance.

NexentaStor supports two types of replication:

- Scheduled replication - Enabled by default with the Enterprise Edition license, with a snapshot schedule of “every 15 minutes” or longer. SR replicates snapshots taken on predefined schedules on the source dataset. If the NexentaStor appliance has the continuous replication license option installed, the snapshot schedule for scheduled replication can be as tight as “every minute”.
- Continuous replication - Requires the continuous replication license option. CR delivers close to-zero Recovery Point Objective (RPO) over any distance without affecting application performance. CR works by asynchronously replicating every write transaction on the source dataset.

To avail of the continuous replication option, contact sales@nexenta.com.

This section lists the CLI commands involved in creating and scheduling replication services. For details on how to prepare your environment for HPR and for advanced replication configurations, see the *High Performance Replication (HPR) Best Practices Guide* in <https://nexenta.com/products/downloads/nexentastor>.

The replication service creation and scheduling process involves the following tasks:

- Set up your replication environment
See [Preparing for HPR](#)
- Check what type of replication feature license is activated in your system
See [Licensing Requirements for the Replication Service](#)
- Optionally, change the HPR password you set up during the NexentaStor installation
See [Managing the HPR Password](#)
- Configure HPR service system attributes.
See [Configuring HPR Address to Send or Receive Replication Traffic](#)
- Create a new replication service.
- List HPR services created.
- Enable/ Disable HPR service.
- On-demand start/ stop a replication service.
- Schedule replication services
See [Configuring a New Replication Service](#)

Definition of Terms

This section uses the following terms:

Table 3-4: Terms and Descriptions

Terms	Definition
Protection service	Protection service is a NexentaStor data and metadata replication service. This replication service is possible only between NexentaStor 5.0 appliances.
Scheduled replication	Scheduled replication is the default type of replication service that comes with the Enterprise Edition license. Scheduled replication generates snapshots of the datasets at the source appliance on a set schedule and replicates these scheduled snapshots on either the local pool on the same NexentaStor appliance or on a remote NexentaStor appliance.
Continuous replication	Continuous replication asynchronously sends every write transaction to the destination dataset and delivers as close to zero Recovery Point Objective as possible over any distance, without affecting performance application.
Source	Source represents the node where the data is synchronized from.
Destination	Destination represents the site where the data is synchronized to and can be either local dataset on the same appliance or a remote NexentaStor appliance. Note: When the replication service is running and if the destination dataset is a filesystem, the filesystem gets unmounted.
Primary	Primary represents the node where the replication service manager is running. See below for more details about Service Manager. So primary appliance can be either source or destination.
Secondary	Secondary represents where the replication service agent is running. See below for more details about Service agent. Secondary appliance can be either source or destination.
Primary appliance (replication manager)	Secondary appliance (replication agent)
Source dataset	destination dataset
Destination dataset	source dataset

Preparing for HPR

For a successful rollout of replication services, the following preparation steps need to be done.

Configure Network

Setup your replication environment.

Configure the logical network:

- Verify the static system management IP you created using the following command. (Note 0.0.0.0 is not supported)


```
CLI@nexenta> config list system.managementAddress
```

- If you need to change the management address after the installation see [Changing Management Address](#)
- Identify available interfaces that will be used to transport the replication traffic.
- Enable jumbo frames for data replication interface to improve replication performance by making data transmissions more efficient.
- Pre-allocate one static IP address for each replication interface.

1. Add a physical network interface card if needed. To add an interface, verify that the new NIC is available on the system.

```
CLI@nexenta> inventory nic
```

NAME	MEDIATYPE	STATE	SPEED	MAC	DUPLEX
e1000g0	Ethernet	up	1000	0:25:90:61:b0:14	full
e1000g1	Ethernet	unknown	0	0:25:90:61:b0:15	half

2. Configure the newly added interface. The following example assigns a static IP address to the e1000g0 NIC.

```
CLI@nexenta> ip create [-ntv] static <name> <address>
```

Example:

```
CLI@nexenta> ip create static e1000g0/v4 10.3.10.38/24
```

3. View the NIC configuration to verify the newly added interface.

```
CLI@nexenta> ip list
```

System response:

NAME	TYPE	STATE	ADDRESS
e1000g0/v4	static	ok	10.3.10.38/24
e1000g1/v6	addrconf	ok	fe80::20c:29ff:fe80:82a6/10

4. If you have already assigned an IP address and need to update the IP address, destroy the IP address and then assign a new IP as in the previous command.

```
CLI@nexenta> ip destroy <name>
```

Example:

```
CLI@nexenta> ip destroy e1000g0/v4
```

- In case of replication from HA to HA, ensure that the management IP on all the 4 appliance nodes are set to static IPs.

Table 3-5: Recommendations for Network Interface Address

Interface	Required and Recommended
Management Network Interface	Required: Fast Ethernet or Gigabit Ethernet adapter, DHCP or static network address. Recommended: dedicated network interface, IPMP or LACP link aggregation.

Interface	Required and Recommended
Replication Data Network Interface	Required: Gigabit Ethernet adapter, static network address. Recommended: dedicated network interface, 10 Gigabit Ethernet, IPMP or LACP link aggregation, jumbo frames.
Cluster Heartbeat Network Interface	Required: Fast Ethernet or Gigabit Ethernet adapter, static network address. Recommended: dedicated network interface.
Cluster Failover Network Interface	Required: Gigabit Ethernet adapter, static network address. Recommended: dedicated network interface, 10 Gigabit Ethernet, IPMP or LACP link aggregation, jumbo frames.

Licensing Requirements for the Replication Service

NexentaStor supports the following two types of replication services:

- Scheduled replication
- Continuous replication

To avail of the continuous replication option, contact sales@nexenta.com.

To verify if the continuous replication service is activated, run the following command. Look for the `continuousReplication` feature in the “features” property.

```
CLI@nexenta> license show
PROPERTY      VALUE
guid          44454c4c-3600-104b-804c-b9c04f4e3232
valid         yes
status        ok
type          ENTERPRISE-TRIAL(Nexenta Internal)
product       NexentaStor
version       5.0.0
licensee      Nexenta-xxxxx@nexenta.com
serial        SR-DEV-NS-201617669
features      allFlash, fibrechannel, highAvailability,
              continuousReplication, scheduledReplication
issued       Thu Sep 29 17:00:00 2016
expires       Sun Nov 13 16:00:00 2016
capacity      no limit
maintenance   Thu Sep 29 17:00:00 2016 - Sun Nov 13 16:00:00 2016 (valid)
```

To activate a license key with the optional continuous replication feature:

```
CLI@nexenta> license activate <key>
```

Managing the HPR Password

When installing the NexentaStor appliance, you must have set the HPR password. Before configuring the HPR service, verify that the replication password is configured and the same on all nodes in the replication group. For security purpose, NexentaStor allows you to configure the HPR service between the nodes that share the same password. If you are unsure that your appliances have the same password, you can reset the password using this command:

```
CLI@nexenta> hpr password-set [--password=<str>]
```

Configuring HPR Address to Send or Receive Replication Traffic

To configure the interface to send or receive all replication traffic on your replication node, see this example:

1. To view existing values for HPR system properties:

```
CLI@nexenta> config list hpr
NAME                                FLAGS  VALUE
hpr.connectTimeout                  --     3000
hpr.dataAddress                      --     -
hpr.dataPort                        --     6000
hpr.heartbeatFaultTolerance         --     3
hpr.heartbeatInterval               --     10000
hpr.requestTimeout                  --     30000
hpr.syncMaxAttempts                 --     12
hpr.syncRetryInterval               --     15000
hpr.totalMemoryLimit                --     25
```

2. To set HPR service data address:

```
CLI@nexenta> config set hpr.dataAddress = <IP address of the node that
will send or receive the replication data>
```

Example:

```
CLI@nexenta> config set hpr.dataAddress=10.3.10.38
```

3. To verify the HPR service data address:

```
CLI@nexenta> config get value hpr.dataAddress
PROPERTY  VALUE
path      hpr.dataAddress
type      variable
value     10.3.10.38
```


Depending on the direction of the replication, the interface that is configured as data address will begin to send or receive replication data.

Note: For successful remote replication, you must configure the HPR data interface explicitly.

Configuring a New Replication Service

The steps below direct you to create a scheduled replication service named `SR-A2B`, between Node A and Node B with the source dataset `poola/primary-fs` residing on Node A and the destination dataset `poolb/secondary-fs` on Node B.

Figure 3-1: Example of a Scheduled Replication Service between two NexentaStor Appliances.



1. Confirm that you are licensed for Scheduled Replication or Continuous Replication.
2. Create a new replication service

You can create a protection service that generates snapshots at the source appliance following a set schedule or on a continuous basis. These snapshots can be stored locally on the source appliance (same appliance where the HPR service was created) or locally and remotely in both the source and destination appliances. In order to create a local-to-remote replication, the remote IP should be added along with the destination dataset. You may also use fully qualified hostnames in the place of IP.

When you select a source dataset for the replication service, you must specify whether to replicate sub-folders of the dataset or include parent dataset in the replication stream. If you want to use the recursive property, you must select the following flag `-r` or `--recursive` in the HPR create command.

```
CLI@nexenta> hpr create [-CeFIlnprv] [--description=<desc>] [--ignore-properties=<names>] [--replace-properties=<property=value>] [--throttle=<value>] [--max-buffer-size=<value>] <type> <source> <destination> <name>
```

The example below creates a local-to-remote replication service named `SR-A2B`, with the remote management IP address specified, with scheduled type (as opposed to `continuous` type), with source `poola/primary-fs` and destination of `poolb/secondary-fs`

```
CLI@nexenta> hpr create -r scheduled poola/primary-fs https://10.3.10.38/24poolb/secondary-fs SR-A2B
```

- Optionally, add a schedule for the service you just created.

The following example configures a schedule called "SR-A2B_sced" for the replication service "SR-A2B" that starts every 15 minutes. This example retains the last 2 snapshots on the source and 3 snapshots on the destination.

```
CLI@nexenta> hpr schedule-add [-nv] <service-name> <cron> <keep-source>
<keep-destination> [<schedule-name>]
```

Example:

```
CLI@nexenta> hpr schedule-add SR-A2B "*/15 * * * *" --keep-source=2 --
keep-destination=3 SR-A2B_sced
```

- Enable the hpr service.

Replication services of type continuous starts immediately once the service is enabled. For scheduled services, replication starts according to the defined schedule.

```
CLI@nexenta> hpr enable <name>
```

Example:

```
CLI@nexenta> hpr enable SR-A2B
```

- Run the service.

To manually trigger the replication service any time, overriding the configured replication schedule, use the following command.

```
CLI@nexenta> hpr start SR-A2B
```

Once the command is run, the HPR service takes snapshots and replicates them from source appliance to destination site/appliance.

- Verify the snapshot was replicated to the destination.

```
CLI@nexenta> hpr get all SR-A2B
```

NAME	PROPERTY	VALUE
SR-A2B	commonSnapshot	hpr-ondemand-2016-07-21-07-21-16-450
SR-A2B	destination	poolb/secondary-fs
SR-A2B	ignoreProperties	-
SR-A2B	isManager	yes
SR-A2B	isRunning	yes
SR-A2B	isSyncing	no
SR-A2B	maxBufferSize	100
SR-A2B	name	rtl_srv1
SR-A2B	recursive	yes
SR-A2B	replaceProperties	
SR-A2B	source	https://smc-53-109:8443/poola/primary-fs
SR-A2B	state	enabled
SR-A2B	type	scheduled

- Check the retention policy and the scheduler.

In any of the supported deployment topologies, you can configure different retention policies at the Primary appliance and at the Secondary appliances.

```
CLI@nexenta> hpr schedules SR-A2B
```

NAME	CRON	KEEPSOURCE	KEEPDESTINATION	DISABLED
SR-A2B_sced	* / 5 * * * * 2		3	no

8. View all the snapshots belonging to the HPR service you just created.

```
CLI@nexenta> hpr snapshots SR-A2B
```

SNAPLISTID	SOURCESNAPSHOTS	DESTINATIONSNAPSHOTS	SERVICE	SCHEDULE
501ca190-8a71-11e6-addd-7d6588903bbb			SR-A2B	SR-A2B_sced
1 * * * * *	2 / 11	2 / 11		

Managing the Created Replication Service

The following table enumerates the basic CLI commands involved in managing the replication services after they are created. For additional information on these commands, type:

- For the man page of the HPR command:

```
CLI@nexenta> man hpr
```

 For its subcommands pre-pended with service.
- To get the list of HPR subcommands:

```
CLI@nexenta> hpr -h
```
- For subcommand usage syntax and options:

```
CLI@nexenta> hpr <subcommand> --help
```

Table 3-6: Managing Created Replication Service

Task	Related CLI Command
List the HPR services created	<pre>CLI@nexenta> hpr list</pre>
Enable/disable an HPR service	Replication services of type continuous starts immediately once the service is enabled. For scheduled services, replication starts according to the defined schedule. <pre>CLI@nexenta> hpr enable <service name></pre> <pre>CLI@nexenta> hpr disable <service name></pre>
Query HPR service attributes	List the property names and values for a specific replication service: <pre>CLI@nexenta> hpr get all <service name></pre>
Modify HPR service attributes	<pre>CLI@nexenta> hpr set <property name>=<value> <service name></pre>
On-demand start/stop of a replication service	To begin or end a replication service, overriding the configured replication schedule, you can use the following subcommands: <pre>CLI@nexenta> hpr start <service name></pre> <pre>CLI@nexenta> hpr stop <service name></pre>

Task	Related CLI Command
Delete a replication service	<pre>CLI@nexenta> hpr destroy [-fnv] [--source-snapshots] [--destination-snapshots] [--destination] <name></pre> <p>The following is an example to force the deletion of a replication service even if it's in a middle of replicating data. A service deletion operation deletes the snapshots created by the service on the source and destination datasets and destroys the destination datasets.</p> <pre>CLI@nexenta> hpr destroy -f --source-snapshots --destination-snapshots --destination SR-A2B</pre>
Clear detected faults	<pre>CLI@nexenta> hpr clear <service name></pre>
Recover broken replication	<pre>CLI@nexenta> hpr recover <service name></pre>

Table 3-7: Managing Snapshots Associated with the Service

Task	Related CLI Command
List snapshots associated with a replication service	<pre>CLI@nexenta> hpr snaplist-find [-s <field>]... [-S <field>]... [-O <flags>] <service name></pre> <p><service name> Name of the replication service</p>
Delete list of snapshots associated with a replication service	<pre>CLI@nexenta> hpr snaplist-delete [-nv] <service name> <snapshotlist-id></pre> <p><service name> Name of the replication service. <snapshotlist-id> Id of snapshot list</p>
Claim snapshots belonging to a deleted schedule	<p>To allow a schedule to claim a list of snapshots that belonged to a deleted schedule, run the following command:</p> <pre>CLI@nexenta> hpr snaplist-claim [-nv] <service name> <schedule name> <snaplist-id></pre> <p>Where the <schedule name> represents the replication schedule that will claim the snapshots.</p>

Scheduling Replication Services

[Table 3-8](#) enumerates the CLI commands involved in managing schedules for replication services. A scheduled replication service replicates data following one or more specified schedule(s).

Table 3-8: Scheduling Replication Services

Task	Related CLI Command
Verify schedule status	CLI@nexenta> hpr schedules <service name>
Enable/disable a schedule	CLI@nexenta> hpr schedule-disable <service name> <schedule name> CLI@nexenta> hpr schedule-enable <service name> <schedule name>
Rename a schedule	CLI@nexenta> hpr schedule-rename <service name> <schedule name>
Delete a schedule	CLI@nexenta> hpr schedule-remove <service name> <schedule name>
Modify a schedule	CLI@nexenta> hpr schedule-set <properties> <Service Name> <schedule name> Disable the HPR service before doing the following actions: <ul style="list-style-type: none"> • modifying its attributes, • adding or removing a schedule, • flipping the direction of replication.

Configuring NDMP for Backups

Prerequisites

In order to back up NexentaStor datasets to a tape drive or to a virtual tape device, you need:

- A backup software compliant with NDMP version 4 installed on a non-NexentaStor appliance. For example, Commvault.
- A tape drive connected to your NexentaStor appliance or a virtual tape drive. NexentaStor enables the NDMP management software to copy NexentaStor datasets to the attached tape drive. Your third party client backup software performs the archiving functions, such as setting the backup schedule.

Configuring and Managing NDMP

You must set up an NDMP service on your NexentaStor appliance in order to enable the client backup software to transfer NexentaStor datasets over the default TCP port 10000.

1. Verify if the NDMP service is online. If not, enable it:

```
CLI@nexenta> svc list ndmp
CLI@nexenta> svc enable ndmp
```

2. To query and set NDMP service attributes, including NDMP version:

```
CLI@nexenta> svc get ndmp
CLI@nexenta> svc set <NDMP property> ndmp
```

Depending on the type of files you want to backup, enable or disable these properties: darSupport, tarBackupFormat, dumpBackupFormat.

3. Enable, disable, and list authentication types (clear-text or cram-md5) using the following commands. The credentials you configure will be used by the backup application to access your NexentaStor datasets.

```
CLI@nexenta> ndmpauth disable [-nv] <auth-type>
CLI@nexenta> ndmpauth enable [-nv] <auth-type> <username> [--
password=pass]
CLI@nexenta> ndmpauth list
```

Troubleshooting

If you are using tape drives, below are several useful commands to use.

To list the tape drives attached to your NexentaStor appliance:

```
CLI@nexenta> inventory tape-device
```

Sample output:

```
CLI@nexenta> NAME PRODUCT REV SERIAL
```

```
CLI@nexenta> /dev/scsi/changer/c2t5000E1112CF4D003d1 3573TL C.30  
00X4U78K2426_LL0
```

To get the the status of a specific tape drive, see example below.

```
CLI@nexenta> mtx -f /dev/scsi/changer/c2t5000E1112CF4D003d1 status
```

Sample output:

```
CLI@nexenta> Storage Changer /dev/scsi/changer/c2t5000E1112CF4D003d1:3  
Drives, 47 Slots ( 0 Import/Export )
```

```
CLI@nexenta> Data Transfer Element 0:Empty
```

```
CLI@nexenta> Data Transfer Element 1:Empty
```

```
CLI@nexenta> Data Transfer Element 2:Empty
```

```
CLI@nexenta> Storage Element 1:Full :VolumeTag=000118L4
```

```
CLI@nexenta> Storage Element 2:Full :VolumeTag=000117L4
```

```
CLI@nexenta> Storage Element 3:Full :VolumeTag=000104L4
```

Using Virus Scan Engines

About NexentaStor vscan

To use NexentaStor vscan service, you must use a third-party VirusScan engine like McAfee on an external host to perform virus scanning operation on files. You may configure multiple scan engines for use by the NexentaStor vscan service. When you use multiple scan engines the file scan requests are distributed among the configured scan engines to balance the load. NexentaStor will issue a scan request to Virus Scan engine every time an open/close operation request is issued from the client to the NAS platform.

This section describes how to use McAfee VirusScan Enterprise for use as a virus scan engine with the NexentaStor 5.0 Appliance vscan service.

Prerequisites

To enable virus scanning services with NexentaStor, the following items need to be in place:

- A 3rd-party virus scanning engine that supports ICAP (for example, McAfee) installed on a server.
- A port available in NexentaStor to communicate with the virus scan engine. The vscan service on the NexentaStor appliance uses port 1344 by default.

Note: Ensure that the port 1344 is not blocked by the firewalls in your environment.

Managing Virus Scan Services on a NexentaStor Appliance

❖ *Set up the vscan service on the NexentaStor 5.0 appliance:*

1. Verify that the virus scan service in your NexentaStor system is online. If not, enable it.

```
CLI@nexenta> svc list vscan
```

```
CLI@nexenta> svc enable vscan
```

2. Query and set virus service attributes.

```
CLI@nexenta> svc get vscan
```

```
CLI@nexenta> svc set <property>=<value> vscan
```

❖ *Set up the vscan engine:*

1. To add, delete, and list virus scan engines.

```
CLI@nexenta> vscan create <vscan engine>
```

```
CLI@nexenta> vscan destroy <vscan engine>
```

```
CLI@nexenta> vscan list
```

2. Query and set virus engine attributes.

```
CLI@nexenta> vscan get <vscan engine>
```

```
CLI@nexenta> vscan set <property>=<value> <vscan engine>
```

- ❖ *Configure vscan on the NexentaStor appliance to point to the McAfee virus scan engines IP address.*

1. After setting up the Virus Scan Engine (VSE), configure vscan on the NXS 5 node to point to the McAfee scanner through our “vscan” command.

```
CLI@nexenta> vscan set host=10.3.53.119 <vscan engine ID>
```

- ❖ *Enable the virus scan option on the filesystem share you want checked for viruses.*

1. Enable vscan on the filesystem.

```
CLI@nexenta> filesystem set vscan=yes poola_fs
```

2. Verify that the vscan was enabled on the filesystem.

```
CLI@nexenta> filesystem get vscan poola_fs
```

Now that you have enabled the vscan service on the filesystem, you can mount the NFS share and create/access data on that share from a different client machine. This will trigger the McAfee scans.

Checking vscan Related Messages in the NexentaStor Syslog

1. To see the most recent 100 syslog messages.

```
CLI@nexenta> journal tail -c 100 messages
```

2. To test if the vscan engine scans for virus, do the following:

1. See <http://www.eicar.org/86-0-Intended-use.html> for suggestions on how to inject a string with a virus that is not harmful.

2. Copy the string on a text file or Word document in a Windows client folder that is mounted on the NexentStor file system you created.

3. Run the following command in the NexentaStor appliance.

```
CLI@nexenta> journal show messages:
```

Virus scanning is successfully running if you see a log entry similar to the message below.

Example of log entry:

```
Sep 23 15:13:23 ns501d vscand: [ID 540744 daemon.notice] quarantine /
poola/pool1/test_virus.txt 1290 - EICAR test file
```

4. To verify from the virus scan engine server, navigate to the VirusScan Console → right click on ICAP AV Scanner → select View Log on the pulldown menu.

Example of log entries:

```
9/23/2016 4:16:22 PM Scan Started WIN-Q2TL2Q9R8OM Scan Request
Received From :ns501d File to scan:
C:\Windows\TEMP\VSEIcapTempFiles\092320161616221_New%20Text%20Document.
txt
```

```
9/23/2016    4:16:22 PM    Scan Result    No Action Taken (Clean failed
because the detection isn't cleanable)    File from filer ns501d scanned
:
C:\Windows\TEMP\VSEIcapTempFiles\092320161616221_New%20Text%20Document.
txt    EICAR test file
```

What Comes Next?

The following chapter covers the details on how to monitor and investigate key aspects of appliance; how to replace a disk; how to send important system information for Nexenta support service (i.e. core dumps, system configuration files, system logs and so on).

Fault Management

This chapter covers the following topics:

- [Viewing NexentaStor Logs](#)
- [Troubleshooting Data Devices](#)
- [Alerts and Events](#)
- [Diagnose Unsuccessful HA Service Start](#)
- [Bundle Services](#)
- [What Comes Next?](#)

This chapter includes pointers in troubleshooting and diagnosing problems you may encounter.

Viewing NexentaStor Logs

The following are useful commands that will help troubleshoot your NexentaStor deployment:

- To view values for the system parameters software version and software list - to query what version of NexentaStor was installed and whether it is activated.

```
CLI@nexenta> config list
```

- To display license terms.

```
CLI@nexenta> license show
```

- To get a summary of system metrics.

```
CLI@nexenta> system status
```

- To list installation and syslog message logs.

```
CLI@nexenta> journal list -o name,file
NAME FILE
install/caiman /var/log/install/install_log
install/nef /var/log/install/nef_log
install/messages /var/log/install/messages
messages /var/adm/messages
```

- To view the system logs

```
CLI@nexenta> journal show messages
```


System response:

```
Oct  8 03:15:00 smc-53-109 krrp: [ID 832813 kern.notice] NOTICE: A new
session h
as been registered (id:[123bc870-8d40-11e6-9050-0153fdec8120])
Oct  8 03:15:00 smc-53-109 krrp: [ID 623046 kern.notice] NOTICE: PDU
Engine conf
ig: dblk_head_sz:[0], dblk_data_sz:[1024], max_mem:[100 MB],
dblk_per_pdu:[513]
, prealloc:[NO]
Oct  8 03:15:00 smc-53-109 krrp: [ID 971085 kern.notice] NOTICE:
Publishing KRRP
event EC_krrp ESC_KRRP_sess_send_done
```

- To see the most recent 100 syslog messages.

```
CLI@nexenta> journal tail -c 100 messages
```

Troubleshooting Data Devices

When a disk is having issues such as hardware errors or bad blocks, you can take the device offline, replace it, and then online the new disk. Use the following steps to perform this task. The following steps are applicable even if the disk is not faulted but you just want to replace the disk.

1. To identify only the failed disks, use the `-x` option in the following command:

```
CLI@nexenta> disk list -x
NAME                               LABEL SIZE MEDIATYPE STATE  LOCATION      ENCLOSURELABEL
USEDNAME
c1t50000393B8C83E9Cd0 - 0.0K hdd      RETIRED 50030480015a6b7f:15 -
```

The disk is failed if the state returns as `RETIRED`.

2. If you need to physically locate the failed disk, use the `blink` feature to make the LED of the failed disk in the JBOD flash.

```
CLI@nexenta> disk indicators --ident=ON c1t50000393B8C83E9Cd0
Setting the ident to ON blinks the LED light for that drive.
```

3. Now verify if the LED is on. The LED is blinking if the `ident` value returns `"Yes"`.

```
CLI@nexenta> disk indicators c1t50000393B8C83E9Cd0
NAME  PROPERTY  VALUE
ident enabled  yes
fail  enabled  no
ok2rm enabled  no
```

4. When you have decided to replace a disk, set the disk in the pool to offline to prevent writing to or reading from the disk. This command is not applicable for spare or cache devices.

```
CLI@nexenta> pool offline <pool> <disk>
```

Example:

```
CLI@nexenta> pool offline poola c1t50000393B8C83E9Cd0
```

5. Replace the disk with another physical disk on the pool. This is equivalent to attaching a new device.

```
CLI@nexenta> pool replace [-fnv] poola c1t50000393B8C83E9Cd0 <new-disk>
```

Example:

```
CLI@nexenta> pool replace [-fnv] poola c1t50000393B8C83E9Cd0
c1t5000CCA39ACF89DAd0
```

The replaced physical device will begin to resilver, detaching the old device. You can then remove the original device from the storage array and it can be used for other purposes.

For optimal performance, Nexenta recommends the new disk to be of the same size, speed, and media type as the existing one.

6. Set the device in the pool to online.

```
CLI@nexenta> pool online [-env] poola c1t5000CCA39ACF89DAd0
```

To reuse the device that is offline in another storage pool, you should first bring the device online in the new pool and then destroy the original storage pool.

Also see the following subcommands:

Task	Related CLI Command
Attaches a new disk to an existing pool disk. The existing disk cannot be part of a raidz configuration. If the disk is not currently part of a mirrored configuration, it automatically transforms into a two-way mirror consisting of the disk and the attached disk. If the disk is part of a two-way mirror, attaching a new disk creates a three-way mirror, and so on. In either case, the new disk begins to resilver immediately.	CLI@nexenta> pool attach
Removes (detaches) a single device from a mirrored vdev from any type of storage pool.	CLI@nexenta> pool detach
Begins a scrub on a given pool. A scrub traverses the data of the entire pool and verifies that all the blocks can be read.	CLI@nexenta> pool start-scrub and pool stop-scrub
Queries and enables drive indicators	CLI@nexenta> disk indicators <diskID>
Indicates that the drive can be removed.	CLI@nexenta> disk indicators -- ok2rm=ON <diskID>
Highlights if the drive is in a failed state.	CLI@nexenta> disk indicators -- fail=ON <diskID>

Alerts and Events

The following commands can be used to troubleshoot problems by providing a detailed list of the system events. These commands list all hardware or software problems, or a transient condition that needs attention (for example, CPU over-utilization). Each problem is represented by a case that is referenced by a unique UUID. You can view the cases and track the associated events. You can also generate low-level reports based on the cases created. Users can modify a case using the repaired, replaced, and acquit subcommands.

1. Show alert cases

```
CLI@nexenta> alert cases [-av] [-u <uuid>] [-c <code>] [-t <time-spec>]
```

Example:

```
CLI@nexenta> alert cases --all
```

UUID	TIME	CODE	SEVERITY
71af9891-2fe9-cbcc-f418-b6771433e700	Sep 22 15:43:09	SMF-8000-YX	major
fd853280-8115-11e6-9483-dd2e65e98d14	Sep 22 15:43:32	NEX-8004	critical

2. List events describing the details associated with a case

```
CLI@nexenta> alert list [-u <uuid>] [-c <code>] [-T <type>] [-t <from>]
```

Example:

```
CLI@nexenta> alert list -u 71af9891-2fe9-cbcc-f418-b6771433e700
```

TIME	EVENTTYPE	CODE	UUID
Sep 22 15:43:09	suspect	SMF-8000-YX	71af9891-2fe9-cbcc-f418-b6771433e700
Sep 27 14:20:32	repaired	FMD-8000-4M	71af9891-2fe9-cbcc-f418-b6771433e700
Sep 27 14:20:33	resolved	FMD-8000-6U	71af9891-2fe9-cbcc-f418-b6771433e700

3. To view the detailed description on the alert case

```
CLI@nexenta> alert cases --verbose --uuid fd853280-8115-11e6-9483-dd2e65e98d14
```

PROPERTY	VALUE
uuid	fd853280-8115-11e6-9483-dd2e65e98d14
time	Sep 22 15:43:32
code	NEX-8004-ED
severity	critical
type	upset
faulty	no
diagnoseEngine	sw://path=/usr/nef/workers/alert/diagLicense.js
description	License for the server has expired
impact	Functionality is restricted to basic operations

```

action          Renew license for the server
response        N/A
suspects

                resource   nef://license
                class      nef.license.time.error
                certainty   100%
                state       repaired

```

4. To view the telemetry report associated with the above alert case:

```
CLI@nexenta> alert reports <case-id>
```

Example:

```

CLI@nexenta> alert reports fd853280-8115-11e6-9483-dd2e65e98d14
PROPERTY  VALUE
reportId  fd849640-8115-11e6-9483-dd2e65e98d14
class     ereport.nef.license.time.error
time      Sep 22 15:43:32
detector  sw://path=/usr/nef/workers/alert/detectLicense.js
resource  nef://license
value     2016-09-22T22:43:32.880Z

```

Also see the following subcommands:

Task	Related CLI Command
Ignore failing case and resolve it.	<pre>alert acquit <uuid></pre> Options: uuid is the identifier of case which should be acquitted.
Notify the system that the faulty resource has been repaired.	<pre>alert repaired <fmri></pre> Options: fmri is identifier of the faulty resource.
Notify the system that the faulty resource has been replaced.	<pre>alert replaced <fmri></pre> Options: fmri is identifier of the replaced resource.

Diagnose Unsuccessful HA Service Start

- Check if the HA feature is enabled on the clustered nodes using the license show command. If not enabled, obtain an HA license token from sales@nexenta.com then run the following command:

```
CLI@nexenta> license activate <key>
```

- Verify if both the nodes resolve to each other by running the following command on both the nodes.

```
CLI@nexenta> net list host

ADDRESS      HOSTNAME    ALIAS    PROTOCOL
::1          localhost  -        ipv6
127.0.0.1    localhost  loghost  ipv4
10.3.53.109  nef06      -        ipv4
10.3.53.110  nef07      -        ipv4
```

- If it is not mapped correctly, update the hosts on both the nodes to ensure that they resolve to each other. On the first node, run the following command.

```
CLI@nexenta> net create host [-nv] <IP address of the second node>
<hostname of the second node>
```

- Run the above command on the second node to update the hosts.
- Verify that the new entry was added correctly using the following command:

```
CLI@nexenta> net list host
```

You should be able to see the IP address of each other.

- Check the state of the HA services. See NexentaStor 5.0 HA QuickStart Guide for the list of possible service states.

```
CLI@nexenta> haservice status - to see status of all HA services
```

```
CLI@nexenta> haservice status <HA service name> - to see status of a
specific HA service
```

Bundle Services

A support bundle is an archive containing important system information for Nexenta support service (i.e. core dumps, system configuration files, system logs and so on). Support bundles can be managed even when the appliance management layer is not functioning, which makes bundles useful for troubleshooting purposes. This is referred to as "fail-safe mode" and it can be enforced for all bundle commands with the offline option.

Creating a Support Bundle

You can create a bundle as a compressed archive on your system. A generated universally unique identifier (UUID) is displayed to reference the created bundle.

```
CLI@nexenta> bundle create
System response:
299c9660-8daf-11e6-8242-69b879618ded
```

Verifying the Bundle

Ensure that the bundling has completed. The following command lists all the supported bundles.

```
CLI@nexenta> bundle list
BUNDLEID                                CREATED          SIZE  ACTION  DONE
299c9660-8daf-11e6-8242-69b879618ded  Oct  8 16:30:59  763.6K  -      -
```

Uploading a Support Bundle

You can upload support bundles to Nexenta's FTP server. This operation may take a long time depending on size of the bundle and your internet connection speed. Both basic and secure FTP protocols are supported.

- ❖ *To query and modify FTP system parameters:*

```
CLI@nexenta> config list support.ftp
CLI@nexenta> config set <support.ftp.property>=value
```

- ❖ *To upload and to cancel the upload of a bundle:*

```
CLI@nexenta> bundle upload 299c9660-8daf-11e6-8242-69b879618ded
CLI@nexenta> bundle cancel 299c9660-8daf-11e6-8242-69b879618ded
```

Upload Status

You may view the status of the upload by using the list command.

```
CLI@nexenta> bundle list
BUNDLEID                                CREATED          SIZE  ACTION  DONE
299c9660-8daf-11e6-8242-69b879618ded  Oct  8 16:30:59  763.6K  upload50%
```

What Comes Next?

NexentaStor supports a variety of advanced storage features. Some of them are discussed in the next chapter.

Advanced Configuration

This chapter covers the following topics:

- [Configuring Write-Back Cache](#)
- [Sharing NFS and SMB Concurrently on a File System](#)
- [Using IPMP](#)

Configuring Write-Back Cache

NexentaStor 5.0 Write Back Cache (WBC) is a new experimental feature that can significantly improve write performance of hybrid pools. As of NexentaStor 5.0, the WBC feature is in experimental state and is not supported in production environments. This restriction will be removed in a future NexentaStor 5 maintenance update.

To use WBC, a pool must first be configured with a Special VDEV comprising a minimum of two mirrored SSDs. Once a pool has a Special VDEV, Write Back Cache can be enabled on file systems, volume groups and volumes in that pool and accelerate write IOs to these datasets. When WBC is enabled on a dataset, all incoming write IOs for that dataset are first handled by the Special VDEV, benefiting from the high IOPS and low latency of the SSDs. The write IOs are then asynchronously flushed to the backend Data VDEV in the pool. This is the main difference between ZIL/SLOG and WBC: in the case of a SLOG, synchronous write IOs are sent to both SLOG and backend Data VDEVs at the same time, while in the case of WBC, all write IOs are first exclusively handled by the Special VDEV, and only later flushed to the Data VDEV, allowing the pool to efficiently pair the random IOPS capability of the Special VDEV SSDs with the throughput capability of the backend Data VDEV. Write IOs effectively flow through the WBC and eventually always end-up on the backend Data VDEVs.

The WBC property is handled like another property and is hierarchical in nature: enabling WBC will apply to all children file systems, or all volumes in a volume group.

Note: Special vdevs should only be of type SSD.
 WBC can be enabled/disabled on a filesystem, if the pool to which it belongs has a Special device configured.
 WBC can be disabled for all pool's datasets and all WBC data blocks can be migrated from special to normal vdev's when removing special device from given pool.

❖ *To create a pool with a special vdev:*

```
CLI@nexenta> pool create [-fnv] [-R altroot] [-o <properties>] <pool>
<vdev>
```

Example:

```
CLI@nexenta> pool create poola mirror 2 c1t5000C50059935FAFd0
c1t5000C500596E4E0Fd0 c1t5000C5005972AB03d0 c1t5000C5005993693Fd0
special mirror c1t5000C50041ABB9BBd0 c1t5000C50041AC5A17d0
```

- ❖ *To view the status of the pool `poola` you created:*

```
CLI@nexenta> pool status <name>
```

Example:

```
CLI@nexenta> pool status poola
pool      poola
health    ONLINE
scan      none requested
trim      none requested
devices
VDEV      HEALTH   READ WRITE CKSUM LOCATION
poola     ONLINE   0   0   0   -
mirror-0  ONLINE   0   0   0   -
c1t5000C50059935FAFd0 ONLINE   0  0  0  5003048000338d3f:5
c1t5000C500596E4E0Fd0 ONLINE   0  0  0  5003048000338d3f:6
mirror-1  ONLINE   0  0  0  -
c1t5000C5005972AB03d0 ONLINE   0  0  0  5003048000338d3f:7
c1t5000C5005993693Fd0 ONLINE   0  0  0  5003048000338d3f:8
special  ONLINE   0  0  0  -
mirror-2  ONLINE   0  0  0  -
c1t5000C50041ABB9BBd0 ONLINE   0  0  0  5003048000311a3f:4
c1t5000C50041AC5A17d0 ONLINE   0  0  0  5003048000311a3f:5
```

- ❖ *To add special vdevs to the created pool:*

```
CLI@nexenta> pool add [-fnv] <pool> <vdev>
```

Example:

```
CLI@nexenta> pool add poola special mirror <list of vdevs to be added to
the pool>
```

- ❖ *To create a file system on `poola`:*

```
CLI@nexenta> filesystem create [-pnv] [-o <properties>] <filesystem>
```

Example:

```
CLI@nexenta> filesystem create poola/smb_test
```

- ❖ *To verify WBC on the created filesystem:*

```
CLI@nexenta> filesystem get wbcache poola/smb_test
```

PATH	PROPERTY	VALUE	SOURCE
poola/smb_test	wbcache	no	INHERITED

❖ *To enable WBC on a vdev:*

```
CLI@nexenta> filesystem set wbcache=yes poola/smb_test
```

You can also disable the write back cache on a dataset subtree by setting the `wbcache` property to `false`, `no`, or `off`.

It is possible to turn off WBC on a dataset, or on all datasets in the pool, let the special VDEV drain to the backend pool and remove the Special VDEV from that pool. The time required to drain an active WBC is typically measured in seconds to minutes.

❖ *Some supported operations:*

- Disable WBC for given dataset (filesystem, volume or volumegroup)
- Remove special vdev from existing pool

❖ *To disable WBC for a given dataset:*

1. Disable by running the following command:

```
CLI@nexenta> filesystem set [-rnv] <properties> <filesystem>
```

Example:

```
CLI@nexenta> filesystem set wbcache=no poola/smb_test
```

2. Wait for data migration to complete.

3. Detach first special mirror component.

```
CLI@nexenta> pool detach [-nv] <pool> <disk>
```

Example:

```
CLI@nexenta> pool detach poola c1t5000C50041AC5A17d0
```

4. Verify if the device is detached.

```
CLI@nexenta> pool status poola
```

```
pool      poola
health    ONLINE
scan      none requested
trim      none requested
devices
VDEV      HEALTH  READ  WRITE CKSUM  LOCATION
poola     ONLINE  0    0    0      -
mirror-0  ONLINE  0    0    0      -
c1t5000C50059935FAFd0 ONLINE  0  0  0  5003048000338d3f:5
c1t5000C500596E4E0Fd0 ONLINE  0  0  0  5003048000338d3f:6
mirror-1  ONLINE  0  0  0  -
c1t5000C5005972AB03d0 ONLINE  0  0  0  5003048000338d3f:7
c1t5000C5005993693Fd0 ONLINE  0  0  0  5003048000338d3f:8
special  ONLINE  0  0  0  -
```

```
c1t5000C50041ABB9BBd0 ONLINE 0 0 0 5003048000311a3f:4
```

5. Remove second special mirror component.

```
CLI@nexenta> pool remove <pool> <disk>
```

Example:

```
CLI@nexenta> pool remove poola c1t5000C50041AC5A17d0
```

Sharing NFS and SMB Concurrently on a File System

While enabling Active Directory users to access NexentaStor NFS and SMB filesystems, 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 is needed. You can use these optional ID mappings 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.

The NexentaStor supports the name-based mapping between Windows and local user IDs and group IDs. Using this option, the IDMAP facility uses the name detail that is stored in the user or group objects in the Active Directory (AD) to map users and groups. For example, an AD object for a particular Windows user or group can be extended to include the corresponding local NexentaStor user or group name.

Creating Name-Based Rules

By default, the name-based mapping is bidirectional. If you use the `-d` option, a unidirectional mapping is created from a Windows user/group to a local one.

To verify that the IDMAP system service is enabled, see the commands below. Optionally, you can enable Identity Management for Unix (IDMU).

```
CLI@nexenta> svc list idmap
```

```
CLI@nexenta> svc set directorybasedmapping=idmu idmap
```

To create a mapping for a Windows group:

```
CLI@nexenta> idmap create -g -name=<name> IDENTITY
```

Options:

<code>-n, --name=<name></code>	Local name
<code>-g, --group</code>	Group mapping
<code>-u, --user</code>	User mapping (default mapping type)

`-d, --unidirectional` Unidirectional mapping

Examples:

```
CLI@nexenta> idmap create -name root tester@MY.DOMAIN
CLI@nexenta> idmap create -name testers -g testerGroup@MY.DOMAIN
```

You can view information about the available name-based mapping rules, including Windows user/group name, local name, identity type and direction.

```
CLI@nexenta> idmap list [<identity>]
```

Using IPMP

Using IP network multipathing (IPMP), you can combine multiple NICs into an IPMP group. These IPMP links can be connected to different switches to increase availability and performance of the network. You can use IPMP on top of link aggregation. IP network multipathing provides load balancing, availability, and performance in the systems with multiple interfaces connecting to one local area network (LAN).

When creating an IPMP interface, provide the primary network interfaces that will be part of the IPMP group and also the standby interfaces that can be used as fallback for the IPMP group to provide network redundancy.

1. Create an IPMP group.

```
CLI@nexenta> ipmp create <IPMP group name>
```

Example:

```
CLI@nexenta> ipmp create group1
```

2. Verify the newly created group.

```
CLI@nexenta> ipmp list
```

System response:

NAME	STATE	LINKS	FDT
group	failed	-	0

3. Add the primary and standby interfaces to the group.

```
CLI@nexenta> ipmp add-member <name> link
```

Example:

```
CLI@nexenta> ipmp add-member group1 e1000g0
```

```
CLI@nexenta> ipmp add-member group1 e1000g1
```

4. Assign an IP address to the created IPMP group.

```
CLI@nexenta> ip create static group1/v4 10.3.10.38
```

5. Create the network route manually.

```
CLI@nexenta> route create [-nv] <destination> <gateway>
```

Example:

```
CLI@nexenta> route create default 10.3.10.1
```

Additional Resources

- In the previous sections, you:
- Configured a network
- Created a pool and added redundancy, cache devices, and logs.
- Created and shared iSCSI volumes, SMB and NFS file systems.
- Configured HA clusters.
- Created, viewed, and scheduled snapshots.
- Created, viewed, and scheduled replication services.

You can apply these steps to your own data center needs augmented by the following resources to extend your knowledge of configuring NexentaStor 5.0 using the CLI. Table below lists the additional documentation that will guide you to manage your storage appliance using REST APIs, the vCenter plugin, and the NexentaFusion GUI. These documents are downloadable from in <https://nexenta.com/products/downloads/nexentastor>.

Companion Resource	Description
NexentaStor 5.0.1	
CLI Reference Guide	Includes the list and descriptions of the NexentaStor CLI commands and UNIX-like utilities available in NexentaStor.
VVOL Admin Guide	Provides policies to manage external storage in virtualized environments.
vCenter Plugin QuickStart	Provides information on how to use the NexentaStor vCenter plugin.
HA QuickStart	Includes CLI steps to configure clusters for high availability (HA).
REST API QuickStart	Shows how to access the REST API documentation online.
High Performance Replication (HPR) Configuration and Best Practices Guide	Lists replication user scenarios and advanced user configuration.
Product Guide	Provides an overview of the NexentaStor capabilities.
Openstack Cinder and Manila driver	Go to docs.openstack.org and search for NexentaStor
Docker volume driver	Go to https://github.com/Nexenta/nexenta-dockerdirver/blob/master/README.md

Companion Resource	Description
NexentaFusion 1.0.0 User Guide	User Guide and GUI online help (click HELP under the COG wheel icon in the top menu bar).
NexentaStor 5.0.1 and NexentaFusion 1.0.0 Installation QuickStart	Provides information on how to deploy and upgrade.

Global Headquarters

451 El Camino Real, Suite 201
Santa Clara, CA 95050
USA

3000-nxs-cliconfig-5.0-000092-A