High Availability Cluster

User Guide

3.1.6

3000-hac-v3.1.6-000047-A

Nexenta reserves the right to make changes to this documentation 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 documentation 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 documentation includes the latest information available at the time of publication.

Product Versions Applicable to this Documentation:

Product	Versions supported
NexentaStor™	3.1.6
HA Cluster	3.1.6

Contents

Preface . . . . **v**

1 Introduction to HA Cluster **1**

About the Nexenta HA Cluster 1

Storage Failover 2

Exclusive Access to Storage 2

About SCSI Reservation 3

Service Failover 3

Sample Network Architecture 3

Additional Resources 4

2 Installing HAC on a New System **5**

About Installation 5

Prerequisites 5

Installing HAC Plugin on a New Cluster 5

- Installing HA Automatically 6
- Installing HA Manually 6

3 Configuring the HA Cluster **9**

About Configuring the HA Cluster 9

Prerequisites 9

Configuring TCP/IP Networking 10

- Resolving Appliance Hostnames 10
- Adding a Virtual IP Address 11
- Modifying the Default Netmasks 11

Binding the Nodes using SSH 12

Configuring the HA Cluster 13

Adding a Shared Volume to the HA Cluster 13

Importing a Shared Volume 15

Removing a Shared Volume 16

4 Heartbeat and Network Interfaces 17
About Heartbeat and Network Interfaces 17
Heartbeat Mechanism 17
Modifying Heartbeat Properties 18

5 Configuring Storage Failover 19
About Configuring Storage Failover 19
Cluster Configuration Data 19
Mapping Information 20
NFS/CIFS Failover 21
Configuring iSCSI Targets for Failover 21
Configuring Fibre Channel Targets for Failover 23
Setting the HA Cluster to ALUA Mode 23
Changing the HBA Port Mode 23
Creating a Target Group 24
Adding WWNs to an Initiator Group 25
Creating a Zvol 26
Mapping a Zvol 26

6 Advanced Setup 27
About Advanced Setup 27
Setting Failover Mode 27
Setting Manual Failover Mode 27
Setting Automatic Failover Mode 28
Adding a Additional Shared Hostnames 28
Enabling the ALUA Mode 30

7 System Operations 31
About System Operations 31
Viewing the HA Cluster Status 31
Manually Triggering a Failover 32
Verifying Shared Volume Status 33
Viewing Support Logs 33

8 Testing and Troubleshooting 35
Repairing a Broken Cluster Service 35
Replacing a Faulted Node 36

9 Upgrading HAC 37

About Upgrading HAC 37

Prerequisites 37

Failing Over the Nodes 37

Upgrading HA Cluster 38

10 Index 39

This page intentionally left blank

Preface

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

Intended Audience

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

Documentation History

The following table lists the released revisions of this documentation.

Table 1: Documentation Revision History

Revision	Date	Description
3000-hac-v3.1.6-000047-A	June, 2014	GA

Contacting Support

Choose a method for contacting support:

- Visit the Nexenta customer portal http://nexenta.force.com/customerportal or partner portal http://nexenta.force.com/partnerportal. Log in and browse a knowledge base.
- Using the NexentaStor user interface, NMV (Nexenta Management View):
  a. Click **Support**.
  b. Select an action:
     - **Send by email**

       Send the support request to the Nexenta support email.
     - **Save to disk**

       Saves the support information to the `/var/tmp` directory on the HA Cluster appliance.
  c. Complete the request form.
  d. Click **Make Request**.

- Using the NexentaStor command line, NMC (Nexenta Management Console):
  a. At the command line, type `support`.
  b. Complete the support wizard.

Comments

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

1

Introduction to HA Cluster

This section includes the following topics:

- About the Nexenta HA Cluster
- Storage Failover
- About SCSI Reservation
- Exclusive Access to Storage
- Service Failover
- Sample Network Architecture
- Additional Resources

About the Nexenta HA Cluster

The Nexenta HA (High-Availability) Cluster consists of two NexentaStor appliances and provides a storage volume-sharing service. Neither system is designated as the primary or secondary system. You manage both systems actively for shared storage, although only one system provides the access to a shared volume at a time. After you create a volume on one server and share it with the other server, then, when HAC detects a system failure, it transfers ownership of the shared volumes to the other server in the Cluster pair.

HA Cluster provides server monitoring and failover. Protection of services, such as iSCSI, involves cooperation with other modules such as the SCSI Target plugin.

An HA Cluster includes:

- **NexentaStor Appliances**

  Runs a defined set of services and monitors each other for failures. HAC connects these NexentaStor appliances through various communication channels, through which they exchange heartbeats that provide information about their states and the services that reside on them.

- **Cluster Service**

  A transferable unit that consists of:

  - Application start-up and shutdown code
  - Network identity and appliance data

    You can migrate services between cluster appliances manually, or automatically, if one appliance fails.

See Also:

- RSF-1 Cluster User Guide

Storage Failover

The primary benefit of HA Cluster is to detect storage system failures and transfer ownership of the shared volumes to the alternate NexentaStor appliance. All configured volume services failover to the other server. HA Cluster ensures service continuity during exceptional events, including power outages, appliances that run out of memory or crash, and other failures.

Currently, the minimum time to detect that an appliance has failed is approximately 10 seconds. The failover and recovery time is largely dependent on the amount of time it takes to failover the data volume on the alternate appliance. Best practices to reduce the failover time include using fewer zvols and file systems for each data volume.

In the default configuration, HA Cluster implements failover storage services if one of the cluster nodes is unavailable. HA Cluster automatically determines which network device to monitor based on the services that are bound to an interface. It checks all nodes in the cluster, so even if a node is not running any services, HA Cluster continues to monitor the unused interfaces. If the state of one changes to offline, it prevents failover to this node for services that are bound to that interface. When the interface recovers, HA Cluster enables failover for that interface again.

Exclusive Access to Storage

You access a shared volume exclusively through the appliance that currently owns the corresponding volume-sharing service. To ensure this exclusivity, HA Cluster provides reliable fencing through the utilization of multiple types of heartbeats. Fencing is the process of isolating a node in an HA Cluster, and/or protecting shared resources when a node malfunctions. Heartbeats, or pinging, allow for constant communication between the servers. The most important of these is the disk heartbeat. Generally, additional heartbeat mechanisms increase reliability of the cluster's fencing logic; the disk heartbeats, however, are essential.

HA Cluster can reboot the failed appliance in certain cases:

- Failure to export the shared volume from the active node to the passive node. This functionality is analogous to Stonith, the technique for fencing in computer clusters.

On disk systems which support a SCSI reservation, you can place a disk before accessing the file systems, and set the system to panic if it loses the reservation. This feature also serves to protect the data on a disk system.

About SCSI Reservation

In a cluster environment multiple hosts have access to the same shared storage. To control access to the shared storage SCSI reservations are used.

HA Cluster supports SCSI-2 reservation. Therefore, when one HA Cluster node has access to a shared volume, it applies SCSI reservation to the disk drives in the shared volume. The reserved disk drives do not process commands from the alternate HA Cluster node until the lock is released.

SCSI reservations are not applied to spare devices or heartbeat drives.

For more information, contact support@nexenta.com.

Service Failover

As discussed previously, system failures result in the failover of ownership of the shared volume to the alternate node. As part of the failover process, HA Cluster migrates the storage services that are associated with the shared volume and restarts the services on the alternate node.

Sample Network Architecture

A sample cluster hardware setup includes:

- Two bare metal x86/64-bit systems with a shared SAS storage
- Two network interface cards (not mandatory, but good practice)

  Nexenta recommends that you separate management and network traffic, so the heartbeat traffic goes through the management network.

The following illustration is an example of an HA Cluster deployment of a Nexenta iSCSI environment. The host server attaches to iSCSI LUNs in the JBOD, which are connected to the Nexenta appliances nodeA and nodeB. The NexentaStor appliances use the active/passive function of the HA Cluster.

The following diagram shows an example of HA Cluster configuration.

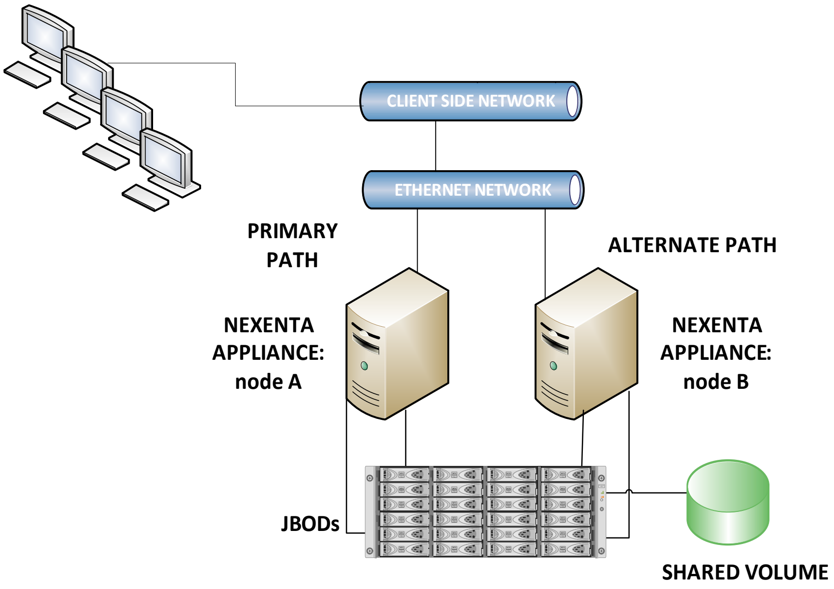

Additional Resources

Nexenta has various professional services offerings to assist with managing HA Cluster. Nexenta strongly encourages a services engagement to plan and install the plugin. Nexenta also offers training courses on high availability and other features.

For more information, contact sales@nexenta.com.

2

Installing HAC on a New System

This section includes the following topics:

- About Installation
- Prerequisites
- Installing HAC Plugin on a New Cluster

About Installation

When installing HAC on a new system, you must install NexentaStor and the HA Cluster plugin software on each NexentaStor appliance in the cluster. This section describes how to set up and install HAC.

Prerequisites

HA Cluster requires shared storage between the NexentaStor Clustered appliances. You must also set up:

- One IP address for each shared volume (zpool)
- Multiple NICs (Ethernet cards) on different subnets for cluster heartbeat and NMV management (this is a good practice, but not mandatory)
- DNS entry for each service name in the cluster

NexentaStor supports the use of a separate device as a transaction log for committed writes. HA Cluster requires that you make the ZFS Intent Log (ZIL) and the L2ARC part of the same storage system as the shared volume.

Installing HAC Plugin on a New Cluster

You can install the HA plugin automatically or manually on a NexentaStor appliance. If you can access your Nexenta repository, you can install HA automatically. If plugin is not available in the repository, contact your Nexenta Sales representative to obtain the HA files.

Installing HA Automatically

You can install the HAC plugin through both NMC and NMV.

❖ *To install a plugin, using NMC:*

1. Verify that plugin is in the repository by typing:

```
nmc:/$ show plugin remotely-available
```

2. Type:

```
nmc:/$ setup plugin install rsf-cluster
```

Note: The plugin is available in the repository after purchasing

3. Confirm the installation.
4. Repeat Step 1 — Step 3 on the other node.

❖ *To install the HAC plugin, using NMV:*

1. Click **Settings > Appliance.**
2. In the Administration panel, click **Plugins**.

Note: The plugins may not be immediately available from your NexentaStor repository. It can take up to six hours before the plugins become available.

3. Click **Add Plugin** for `rsf-cluster` in the Remotely-available plugins section.
4. Confirm the installation.
5. Repeat Step 1 — Step 4 on the other node.

Installing HA Manually

Certain scenarios may require you to install the HA packages manually. For example, if you are installing HA on the appliance that does not have an internet connection.

❖ *To install HA manually:*

1. Obtain the installation files from your Nexenta Sales representative.
2. Switch to a BASH shell. Type:

```
# option expert_mode=1

# !bash
```

3. Type Y to confirm the switch to bash.
4. Use WinSCP, or other Secure Copy Protocol application, to transfer the HA `*.deb` files to the `/root` directory on your NexentaStor appliance.
5. Type the following to unpack and install the HA files:

```
# dpkg --force-all -i *.*
```

6. Type the following to change back to NMC:

```
# nmc
```

7. Type the following to restart NMS:

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

8. Repeat Step 1 — Step 7 on the other node.

This page intentionally left blank

3

Configuring the HA Cluster

This section includes the following topics:

- About Configuring the HA Cluster
- Prerequisites
- Binding the Nodes using SSH
- Configuring the HA Cluster
- Configuring the HA Cluster
- Adding a Shared Volume to the HA Cluster
- Importing a Shared Volume
- Removing a Shared Volume

About Configuring the HA Cluster

You can configure and manage the HA Cluster through the appliance’s web interface, the Nexenta Management View (NMV), or the Nexenta Management Console (NMC).

Note:	This section applies to new installations of HA Cluster. When upgrading, you save and restore the configuration from your previous cluster, so the following sections do not apply.

Prerequisites

Before configuring the HA Cluster, complete the following tasks:

- Configure two bare metal NexentaStor appliances.

  See *NexentaStor Installation Guide*.
- Connect a shared storage to the NexentaStor appliances.
- Create a shared volume.

  See *NexentaStor User Guide*, section *“Creating a volume”*.

- Configure virtual IP address for the shared volume and modify the NexentaStor host table.

  See Configuring TCP/IP Networking.
- Bind the NexentaStor appliances using SSH.

  See Binding the Nodes using SSH.

Configuring TCP/IP Networking

To configure HA Cluster and resolve the HA Cluster nodes hostnames, you must have correct entries in your `/etc/hosts` file. The `/etc/hosts` file on the NexentaStor appliance is a simple text file that associates IP addresses with hostnames.

See Also:

- Adding a Virtual IP Address
- Modifying the Default Netmasks

Resolving Appliance Hostnames

The appliances in the HA Cluster group must be resolvable to each other. This means they must be able to detect each other on the network and communicate. To achieve that, add hostname - virtual IP address pairs into the NexentaStor Internet host table file.

For each host a single line should be present with the following information:

```
IP_address   hostname   [aliases...]
```

❖ *To resolve appliance hostnames:*

1. Log in to NMC on one of the NexentaStor appliances.
2. Type the following to open the `/etc/hosts` file:

   ```
   nmc:/$ setup appliance hosts
   ```
3. Using the vim editor commands, modify the appliance hostnames - IP addresses pairs.

   You must add the hostname - IP address pair of the other HA Cluster node, as well as verify that the local node is configured correctly.

   Example:

   ```
   Internet host table
   ::1   localhost
   127.0.0.1   localhost
   192.168.11.1  <nodeA nodeA.example.com> loghost
   192.168.10.1  <nodeB nodeB.example.com>
   ```
4. Save the changes by typing:

   ```
   :!x
   ```
5. Log in to other node.

6. Repeat Step 2 - Step 4.

Adding a Virtual IP Address

There is a name associated with a shared volume service that is referred to as a **virtual shared service hostname,** or **virtual IP address (VIP)**. The network clients use the virtual hostname to connect to the shared volume.

You must specify a VIP and a corresponding virtual shared service hostname in the `/etc/hosts` file.

For each host a single line should be present with the following information:

```
IP_address   hostname   [aliases...]
```

❖ *To verify hostnames, using NMC:*

1. Log in to NMC on one of the NexentaStor appliances.
2. Type the following to open the `/etc/hosts` file:

   ```
   nmc:/$ setup appliance hosts
   ```

3. Using the `vim` editor commands, type a virtual IP address for the shared volume.

   Example:

   ```
   Internet host table
   ::1   localhost
   127.0.0.1   localhost
   192.168.11.1  <nodeA nodeA.example.com> loghost
   192.168.10.1  <nodeB nodeB.example.com>
   192.168.1.1  <shared_hostname>
   ```

Note: Use the failover hostname to add a shared volume.

4. Repeat Step 1 — Step 3 for other node.

See Also:

- Adding a Shared Volume to the HA Cluster

Modifying the Default Netmasks

When you add a volume service to the HA Cluster, NexentaStor assigns a default netmask for the class of IP network it is on. However, you may need to change the default netmasks according to your configuration changes.

❖ *To modify the default netmask, using NMC:*

1. Log in to one of the NexentaStor appliances.
2. Type the following to open the `/etc/netmasks` file:

   ```
   nmc:/$ setup appliance netmasks
   ```

3. Add the netmask for each network address:

   Example:

   ```
   192.168.1.0         255.255.255.0
   192.168.13.0        255.255.255.0
   192.168.0.0         255.255.0.0
   ```

4. Repeat Step 1 — Step 3 for other node.

Binding the Nodes using SSH

Before you configure the SSH bindings, complete the steps in Resolving Appliance Hostnames. You must bind the two HA nodes together using the SSH protocol so that they can communicate.

Note: You must bind the appliances using either default failover hostnames or the hostnames that you specified in the `/etc/hosts` file. HA Cluster does not support binding using IP addresses.

❖ *To bind the two nodes, using NMV:*

1. Click **Settings > Network**.
2. In the **Network** panel, click **SSH-Bind**.
3. Type the following for the remote server:
   - Server name
   - User Name
   - Password
4. Click **Bind**.

❖ *To bind the two nodes, using NMC:*

1. Log in to one of the NexentaStor appliances.
2. Type the following:

   ```
   nmc:/$ setup network ssh-bind
   ```

3. Type the host name of the NexentaStor appliance that you want to bind.

Note: You must have corresponding entries in the `/etc/hosts` file on both HA Cluster nodes.

For mode information, see

4. Type the root password.
5. Repeat Step 1 — Step 4 on the other node.
6. To verify that you set up the bindings correctly, type:

   ```
   nmc:/$ show network ssh-bindings
   ```

Configuring the HA Cluster

Before you configure the HA Cluster, verify that you completed the steps in Adding a Virtual IP Address and Binding the Nodes using SSH. You need to configure multiple options for HA Cluster before you can use it successfully.

Note: You cannot assign a NexentaStor appliance to more than one HA Cluster.

❖ *To configure an HA cluster, using NMV:*

1. On the first node, select **Settings > HA Cluster**.
2. In the Cluster Settings panel, click **Initialize.**
3. Type or change the **Cluster name**.
4. Optionally, type a description.

Note: If you bound the appliances using SSH, the hostnames of NexentaStors are automatically selected as nodes of the HA Cluster.

5. Select **Enable Network Monitoring**

   The cluster monitors network for nodes.
6. Click **Configure**.
7. Click **Yes**.

❖ *To configure an HA cluster, using NMC:*

1. Type:

   ```
   nmc:/$ create group rsf-cluster
   ```
2. Follow the on-screen instructions.
3. Verify that you created the group, type:

   ```
   nmc:/$ show group rsf-cluster
   ```

Adding a Shared Volume to the HA Cluster

After you configure an HA Cluster, you must add a shared volume to the volume service. You can add a shared volume to the Cluster any time.

❖ *To add a volume to an HA Cluster, using NMV:*

1. In the Cluster Settings panel, click **Volumes**.
2. Select a volume from the drop-down menu.
3. Type the **Failover Hostname** and netmask.

   Use the failover hostname that you have previously specified in the `/etc/hosts` file.

Example:

```
192.168.60.22        <shared_hostname>
```

Note: You can also type *NONE*.

4. Select **Primary appliance**.
5. Optionally, modify initial and standard timeouts for the HA Cluster services.
6. Select **Heartbeat** devices.

   If you use multiple JBODs, select drives from different JBODs as heartbeat devices. Good practice is to assign 2 — 6(max) heartbeat drives.

Warning: Do not assign SSDs as heartbeat devices. Heartbeat devices perform a constant periodic write of the heartbeat signature from each node which can reduce the life of SSDs.

7. Select failover interfaces for both nodes.

Note: To ensure better performance and reliability, use non-primary network interfaces for heartbeats.

8. Click **Add this volume to the cluster**.
9. Click **Confirm**.

❖ *To create a shared service, using NMC:*

1. Type:

   ```
   nmc:/$ setup group rsf-cluster <cluster_group_name>
   shared-volume add
   ```

   System response:

   ```
   Shared volume                  : <shared_volume>
   ```

2. Select heartbeat devices.

   System response:

   ```
   HB disks for 'shared_volume'          : c1t5d0, c1t6d0
   ```

3. Specify a virtual IP address for the selected volume.

Example:

```
VIP1 Shared logical hostname   : <shared_hostname>
```

Note: If you did not define the hostname/IP address for the shared volume, NMC prompts you to modify the NexentaStor host table.

Example:

```
Internet host table
::1   localhost
127.0.0.1   localhost
192.168.60.107  <nodeA nodeA.example.com> loghost
192.168.60.79  <nodeB nodeB.example.com>
192.168.60.22  <shared_hostname>
```

4. Select network interfaces for the selected VIP(s) on both HA Cluster nodes.
5. Specify **Failover Netmask**.

   Example:

   ```
   VIP1 Failover Netmask         : 255.255.255.0
   ```

   System response:

   ```
   Stop adding VIPs?  (y/n)
   ```

6. Confirm the VIP configuration by typing `y`.
7. If you want to add additional VIPs, type `n` and repeat Step 3 - Step 5.
8. Select the `Main node`.

   Main node or Primary node is the NexentaStor appliance where the shared volume is imported after you complete the HA Cluster configuration.

9. Optionally, modify the `Initial timeout` and `Standard timeout`.

   System response:

```
Enable SCSI PGR reservation by typing y.
Adding shared volume '<shared-volume>', please wait ...
Jan  8 23:20:49 nodeA RSF-1[19746]: [ID 702911 local0.alert] RSF-1 cold restart:
All services stopped.
Waiting for add operation to complete ........... done.

HA CLUSTER STATUS: HA-Cluster
NodeA:
 <shared-volume> running      auto   unblocked  <shared_hostname e1000g1   20  8
NodeB:
 <shared-volume>  stopped       auto    unblocked  <shared_hostname> e1000g1
20  8
```

Importing a Shared Volume

To put a shared volume under the cluster control, you may need to import it. The volume may get exported when you upgrade a NexentaStor appliance or if a NexentaStor appliance is unavailable.

❖ *To import a shared volume, using NMV:*

1. Click **Data Management > Data Sets.**
2. In the **All Volumes** panel, click **Import**.
3. Click **Import** for the relevant volume.

❖ *To import a shared volume, using NMC:*

◆ Type:

```
nmc:/$ setup volume import
```

Removing a Shared Volume

You can remove a shared volume from the HA Cluster control. The remove operation does not delete any data on the volume. The volume remains imported on the active HA Cluster node. However, the volume becomes unavailable through the failover hostname.

❖ *To remove a shared volume, using NMV:*

1. In the Cluster Settings panel, click **Volumes**.
2. Select the **Remove a volume** tab.
3. Select a volume.
4. Click ❌.

❖ *To remove a shared volume, using NMC:*

1. Type:

```
nmc:/$ setup group rsf-cluster <cluster_group_name>
shared-volume <shared_volume> remove
```

System response:

```
Remove shared volume <shared_volume> and restart HA
Cluster <cluster> ?   (y/n)
```

2. Confirm the operation by typing **y.**

System response:

```
Removing shared volume <shared_volume>, please wait ...
```

4

Heartbeat and Network Interfaces

This section includes the following topics:

- About Heartbeat and Network Interfaces
- Heartbeat Mechanism
- Modifying Heartbeat Properties

About Heartbeat and Network Interfaces

A NexentaStor appliance in the HA Cluster constantly monitors the state and status of the other appliance in the cluster through heartbeats. Because HA Cluster servers must determine that an appliance (member of the cluster) has failed before taking over its services, you configure the cluster to use several communication channels through which to exchange heartbeats.

Heartbeat Mechanism

The loss of all heartbeat channels represents a failure. If an appliance wrongly detects a failure, it may attempt to start a service that is already running on another server, leading to so-called *split brain* syndrome. This can result in confusion and data corruption. Multiple, redundant heartbeats prevent this from occurring.

HA Cluster supports the following types of heartbeat communication channels:

- **Heartbeat/Reserved devices**

  Accessible and writable from all appliances in the cluster or VDEV labels of the devices in the shared volume.

  If you select the heartbeat devices, VDEV labels for devices in the shared volume perform the heartbeat function. If a shared volume consists of a few disks, NexentaStor uses VDEV labels for one or more disks for the heartbeat mechanism. You can specify which disks.

The heartbeat mechanism uses sectors 512 and 518 in the blank 8K space of the VDEV label on each of the selected heartbeat devices. Therefore, the heartbeat drives are still used to store data.

- **Network Interfaces**

  The preferred heartbeat connection is a dedicated "cross-over" connection between the nodes. You can also use or add any other interface type (simple, IPMP, aggregate) for additional resiliency.

- **Serial Links**

  Heartbeat connection using a null-modem serial cable plugged into both NexentaStor nodes.

Note: Nexenta recommends to use network interfaces and reserved devices for cluster heartbeat.

Modifying Heartbeat Properties

When you configure a cluster, you define the heartbeat properties. However, you can modify them later, if needed.

This functionality is only available in NMC.

Note: You can include a NexentaStor only to one HA Cluster.

❖ *To change heartbeat properties, using NMC:*

1. Type:

   ```
   nmc:/$ setup group rsf-cluster <cluster_name>
   hb_properties
   ```

   System response:

   - Enable inter-appliance heartbeat through primary interfaces?: **Yes | No**
   - Enable inter-appliance heartbeat through serial ports?: **Yes | No**
   - Proceed: **Yes | No**

2. Follow the on-screen instructions.

5

Configuring Storage Failover

This section includes the following topics:

- About Configuring Storage Failover
- Cluster Configuration Data
- Mapping Information
- NFS/CIFS Failover
- Configuring iSCSI Targets for Failover
- Configuring Fibre Channel Targets for Failover

About Configuring Storage Failover

HA Cluster detects storage system failures and transfers ownership of shared volumes to the alternate NexentaStor appliance. HA Cluster ensures service continuity in the presence of service level exceptional events, including power outage, appliance running out of memory, or crashing, etc.

The HA Cluster failover does not repair disk failures. Therefore, if one or more disks in the shared volume are unreachable, you cannot repair them using failover. Moreover, HA Cluster does not detect disk and volume failures. If you execute failover for a degraded volume, the failover operation may take long time or hang the system. Therefore, Nexenta recommends that you always verify that shared volume is in healthy state, before executing failover.

Cluster Configuration Data

When you configure SCSI targets for either FC or iSCSI in a cluster environment, make sure that you are consistent with configurations and mappings across the cluster members. HA Cluster automatically propagates all SCSI Target operations. However, if the alternate node is not available or not configured at the time of the configuration change, problems can occur. By default, the operation results in a warning to the User that the remote update failed.

Execute all FC and iSCSI configuration changes after you place the volume under cluster control. Otherwise, the changes will not be reflected on the other node.

You can also set HA Cluster to synchronous mode. In this case, the action fails completely if the remote update fails.

❖ *To set the synchronous mode, using NMC:*

1. Type:

   ```
   nmc:/$ setup appliance nms property
   rsf_config_update_synchronous
   ```

   System response:

   ```
   View or modify NMS property
   'rsf_config_update_synchronous'. RSF-1 Appliance
   configuration update mode. 1 - Strict Synchronous update
   across the cluster, 0 - Asynchronous update, if synchronous
   update fails. Navigate with arrow keys (or hjkl), or Ctrl-
   C to exit.
   ```

2. Select an appropriate value:

   - **1** — Strict Synchronous update across the cluster
   - **0** — Asynchronous update if synchronous update fails.

To protect local configuration information that did not migrate, periodically save this configuration to a remote site (perhaps the alternate node) and then use NMC commands to restore it in the event of a failover.

The `restore` command restores previously saved configuration data that includes:

- Target groups
- Host groups (`stmf.config`)
- Targets
- Initiators
- Target portal groups (`iscsi.conf`)

Mapping Information

Use SCSI Target to map zvols from the cluster nodes to client systems. It is critical that the cluster nodes contain the same mapping information. Mapping information is specific to the volume and is stored with the volume itself.

The cluster software attempts to keep the HA configuration in sync between the nodes. In certain situations, such as making SCSI Target changes with one of the two cluster nodes offline, the configuration may become out-of-sync. An out-of-sync configuration takes longer to fail over and in some cases may require administrator's actions to resolve.

NFS/CIFS Failover

You can use HA Cluster to ensure the availability of NFS shares to users. However, note that HA Cluster does not detect the failure of the NFS server software.

NFS/CIFS settings are volume-level properties that migrate between nodes automatically upon failover. However, settings such as idmap may need to be defined on both nodes.

HA Cluster does not detect CIFS server failures.

Configuring iSCSI Targets for Failover

You can use HA Cluster to failover iSCSI volumes from one cluster node to another. The target IQN moves as part of the failover.

Setting up iSCSI failover involves setting up a zvol in the shared volume.

Note: You create and share a zvol through iSCSI separately from the HA Cluster configuration.

If you create iSCSI zvols before marking the zvol's volume as a shared cluster volume, then when you share the cluster volume as an active iSCSI session, it may experience some delays. Depending on the network, application environment and active workload, you may also see command level failures or disconnects during this period.

When you add a shared volume to a cluster which has zvols created as back up storage for iSCSI targets, it is vital that you configure all client iSCSI initiators, regardless of the operating system, to access those targets using the shared logical hostname that is specified when the volume service was created, rather than a real hostname associated with one of the appliances.

The cluster manages all aspects of the shared logical hostname configuration. Therefore, do not configure the shared logical hostname manually. Furthermore, unless the shared volume service is running, the shared logical hostname is not present on the network, however, you can verify it with the ICMP ping command.

❖ *To configure iSCSI targets on the active appliance, using NMV:*

1. Click **Data Management > SCSI Target.**
2. In the zvols panel, click **create**.
3. Make the virtual block device > 200MB.

   HAC automatically migrates the newly created zvol to the other appliance on failover. Therefore, you do not have to duplicate the procedure manually.

4. From the iSCSI pane, click **iSCSI > Target Portal Groups** and define a target portal group.

⚠ **Note:** It is critical that the IPv4 portal address is the shared logical hostname specified when the volume service was created, instead of a real hostname associated with one of the appliances.

HAC automatically replicates the newly created target portal group to the other appliance.

5. Create an iSCSI target and add it to the target portal group:
   a. Click **iSCSI > Targets**.
   b. Type a name and an alias.
   c. Select the Target Portal Group that you created in Step 4.

      The newly created iSCSI target displays in the Targets page.

      A target portal group limits zvol visibility to remote client initiators. The newly created iSCSI target is automatically replicated to the other appliance.

6. Create a LUN mapping to the zvol, using NMV:
   a. From the SCSI Target pane, click **Mappings**.
   b. Select a zvol.
   c. Optionally, select initiator group, target group, and LUN number.

      You can manage the visibility of the zvol by mapping it to different target and portal groups.

      For more information, see section “Managing SCSI Targets” in the *NexentaStor User Guide*.

      The operation creates a LUN mapping to the zvol for use as backup storage for the iSCSI target. The newly created LUN mapping is automatically migrated to the other appliance on failover.
   d. On the client, configure the iSCSI initiator to use both the IQN of the iSCSI target created and the shared logical hostname associated with both the volume service and the target portal group to access the zvol through iSCSI.

Failover time varies depending on the environment. As an example, initiating failover for a pool that contains six zvols, the observed failover time is 32 seconds. Nodes may stall while the failover occurs, but otherwise recover quickly.

See Also:

- “Managing SCSI Targets” in the *NexentaStor User Guide*
- *SCSI Target FC User Guide*

Configuring Fibre Channel Targets for Failover

To configure the fiber channel targets for HA Cluster failover, you must complete the following tasks:

- Setting the HA Cluster to ALUA Mode
- Changing the HBA Port Mode
- Creating a Target Group
- Adding WWNs to an Initiator Group
- Creating a Zvol
- Mapping a Zvol

Setting the HA Cluster to ALUA Mode

The FC target is tied to the WWN. WWN is hardware based and since each FC HBA has a different HW WWN, it cannot be fallen over. When you enable the Asymmetric Logical Unit Access (ALUA) mode, NexentaStor nodes can communicate the FC ports on the second node with the first.

Therefore, COMSTAR presents the exported LUNs on the second node in `STANDBY` state. This would be similar to having the same LUN presented via two different IP addresses for iSCSI. Normal running mode is to have the first head doing the FC work with the FC port with the LUNS in `ACTIVE` mode on it and the second NexentaStor node would have the same LUNs in `STANDBY` mode. In case of a failover, the `ACTIVE` port becomes `STANDBY` and the `STANDBY` becomes `ACTIVE`.

❖ *To set the HA Cluster to ALUA mode:*

1. Log in to an HA Cluster node.
2. Click **Settings > HA Cluster**.
3. Select Advanced > **Miscellaneous Options**.
4. Select the **Enable ALUA mode** checkbox.

See Also:

- Adding WWNs to an Initiator Group
- Creating a Zvol
- Creating a Target Group

Changing the HBA Port Mode

As a prerequisite for configuring Fibre Channel targets, change the HBA port modes of both appliances from Initiator mode to Target mode.

❖ *To change the HBA port mode, using NMV:*

1. Click **Data Management > SCSI Target Plus**

2. **Select Fibre Channel > Ports.**
3. Select **Target** from the Mode dropdown menu.
4. Once you change the HBA port modes of both appliances from Initiator mode to Target mode, reboot both appliances so the Target mode changes can take effect.

See Also:

- Creating a Target Group
- Setting the HA Cluster to ALUA Mode
- Creating a Zvol

Creating a Target Group

To use ALUA, create an FC target using the desired FC ports from both nodes. Failure to create a target group with FC ports from both nodes may result in the inability of a SCSI client to maintain access to the storage upon a failover event.

Target FC ports are identified as:

- **Local**

  An FC port that resides on the node that you are configuring.

- **Remote**

  An FC port that resides on the second node in the same cluster.

To ensure proper failover, a Target Group must have at least one local and one remote port defined within it.

❖ *To create a target group, using NMV:*

1. Click **Data Management > SCSI Target Plus**.
2. In the SCSI Target panel, click **Target groups**.
3. Click **Create** or **here**.
4. In the **Group Name** field, type the name of the target group.
5. Select at least one local and one remote FC ports.

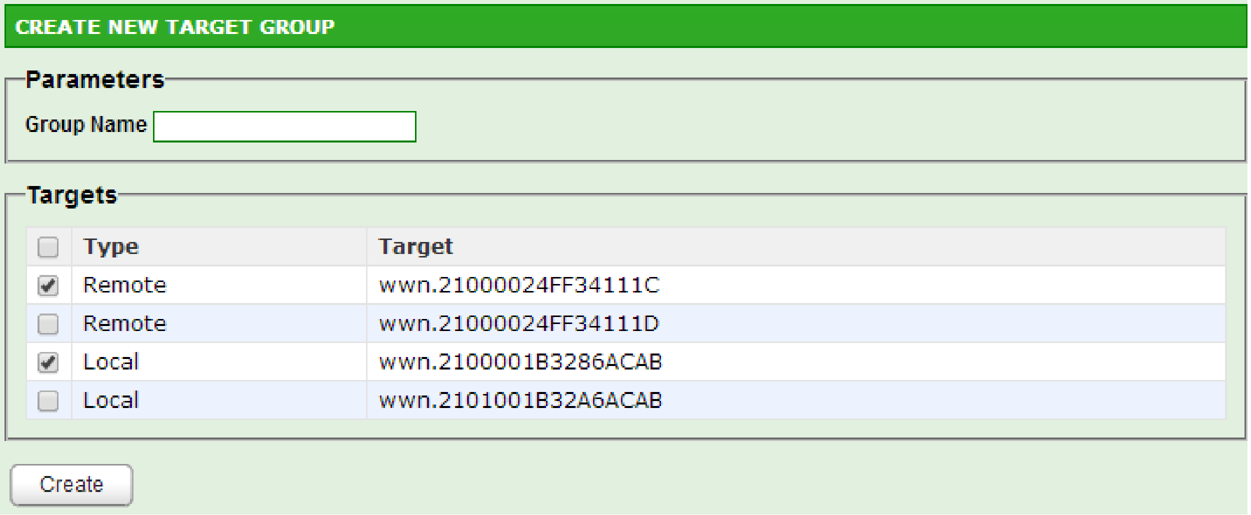

CREATE NEW TARGET GROUP

Parameters

Group Name

Targets

☐	Type	Target
☑	Remote	wwn.21000024FF34111C
☐	Remote	wwn.21000024FF34111D
☑	Local	wwn.2100001B3286ACAB
☐	Local	wwn.2101001B32A6ACAB

See Also:

- Adding WWNs to an Initiator Group
- Creating a Zvol
- Mapping a Zvol

Adding WWNs to an Initiator Group

You must configure an initiator group for each SCSI client, or group of clients. Depending on the configuration of your SAN, some SCSI client initiators may not be visible on some nodes. Therefore, you may need to manually add the client initiator information for initiators not visible to the node that you are working on. Alternatively, you can create an initiator group on one node and assign initiators that are visible to that node. Then you need to modify the initiator group from the other node by adding initiators visible from that second node.

❖ *To create an initiator group, using NMV:*

1. Log in to an HA Cluster node.
1. Click **Data Management > SCSI Target**.
2. In the **SCSI Target** panel, click **Initiator Groups**.
3. In the **Manage Groups of Remote Initiators** window, click **here**.
4. In the **Create New Initiator Group** window:
5. In the field **Group Name**, specify a custom group name.
6. In the **Additional Initiators** field, type the WWNs of additional initiators, not visible to this node, separated by comma.
7. Click **Create**.

See Also:

- Creating a Zvol
- Mapping a Zvol
- Creating a Target Group

Creating a Zvol

Create a zvol using the required FC LUNs.

❖ *To create a zvol, using NMV:*

1. Click **Data Management > SCSI Target**.
2. In the **ZVOLS** panel, click **Create**.
3. In the **Create a New ZVOL** window, fill in the required fields and click **Create**.
4. Proceed to Mapping a Zvol.

   For more information, see *NexentaStor User Guide*.

See Also:

- Mapping a Zvol
- Adding WWNs to an Initiator Group
- Creating a Target Group

Mapping a Zvol

Map the zvol that you created using the FC LUNs to appropriate initiator and target groups to ensure LUN visibility and failover capability.

❖ *To map a zvol:*

1. Click **Data Management > SCSI Target**.
2. In the **ZVOLS** panel, click **Mapping**.
3. In the **Manage Mappings** window, click **here**.
4. In the **Create New Mapping** dialog, fill the required fields.
5. Click **Create**.

See Also:

- Creating a Zvol
- Creating a Target Group
- Adding WWNs to an Initiator Group

6

Advanced Setup

This section includes the following topics:

- About Advanced Setup
- Setting Failover Mode
- Adding a Additional Shared Hostnames
- Enabling the ALUA Mode

About Advanced Setup

This section describes advanced functions of HA Cluster, such as setting the failover mode, adding virtual hostnames and volumes, and other miscellaneous options.

Setting Failover Mode

The failover mode defines whether or not an appliance attempts to start a service when it is not running. There are separate failover mode settings for each appliance that can run a service.

Note: Set failover mode to automatic every time you perform any maintenance to avoid unwanted failover events.

You can set the failover to the following modes:

- Setting Manual Failover Mode
- Setting Automatic Failover Mode

Setting Manual Failover Mode

In manual mode, the HA Cluster service does not initiate the failover when it detects a failure. However, it generates warnings when the parallel appliance is not available. If the appliance cannot obtain a definitive answer about the state of the service, or the service is not running anywhere else, the appropriate timeout must expire before you can take any action. The primary service failover modes are typically set to automatic to ensure that an appliance starts its primary service(s) on boot up.

⚠ **Note:** Setting a service to manual mode when the service is already running does not stop that service, it only prevents the service from starting on that appliance.

❖ *To set the failover mode to manual, using NMV:*

1. Click **Advanced Setup > Cluster Operations > Set all Manual**.
2. Click **Yes** to confirm.

⚠ **Note:** Before HAC performs an operation, it saves the state of the services in the cluster, which you can later re-apply to the cluster using the restore button. Once HAC restores the service state, HAC clears the saved state.

❖ *To set the failover mode to manual, using NMC:*

◆ Type:

```
nmc:/$ setup group rsf-cluster <cluster_name> shared-
volume <volume_name> manual
```

Setting Automatic Failover Mode

In automatic mode, the appliance attempts to start the service when it detects that there is no available parallel appliance running in the cluster. Automatic failover mode is the default setting.

❖ *To set the failover mode to automatic, using NMV:*

1. Click **Advanced Setup > Cluster Operations > Set all Automatic.**
2. Click **Yes** to confirm.

❖ *To set the failover mode to automatic, using NMC:*

◆ Type:

```
nmc:/$ setup group rsf-cluster <cluster_name> shared-
volume <volume_name> automatic
```

❖ *To stop all services in the Cluster, using NMV:*

1. Click **Stop All Services**.
2. Click **Yes** to confirm.

Adding a Additional Shared Hostnames

You add a VIP, or shared hostname, which network clients use to access the shared storage when you create an HA Cluster. You can also add more than one shared hostname later. Additional shared hostnames provide the access to a shared volume through alternate IP addresses.

❖ *To add a virtual IP address, using NMV:*

1. In the **Cluster Settings** panel, click **Advanced**.
2. Click **Additional Virtual Hostnames**.
3. Select a shared volume from the drop-down list.
4. Click **Add a new virtual hostname**.Type the virtual hostname and netmask.
5. Select an interface for each node.
6. Click **Add**.
7. If prompted, type the IP address of the failover node.

Note: Use the IP address that is not in use and that is accessible from both nodes of the HA Cluster. You can add the hostname and IP address pairs to the NexentaStor host table.

See Adding a Virtual IP Address.

8. Click **Add**.
9. Click **Save Settings.**
10. Click **OK** to confirm the modifications.

❖ *To add a virtual IP address, using NMC:*

1. Type:

```
nmc:/$ setup group rsf-cluster <HA Cluster> vips add
```

2. Select the HA Cluster service.
3. Type a virtual hostname.
4. If you type the IP address or hostname, that one or more HA Cluster nodes cannot resolve, NexentaStor prompts you to modify the local host tables.
   - If you want to modify the local host tables:
     1) Type **y**.
     2) Type the IP address and host name.
   - Alternatively, you can configure the DNS server settings.
     1) Type **n**.
     2) Log in to your DNS server and add the host name and IP address pair to the DNS settings.
     3) Repeat Step 1 — Step 3 and Step 5 — Step 10.
5. Select a network interface for this node.

   Nexenta recommends that you configure additional network interfaces rather than specifying the primary network interface.
6. Select network interface for the remote node.
7. Type the failover netmask.
8. Confirm the settings by typing **y**.

System response:

```
Stop adding VIPs?  (y/n)
```

9. Type **y** to finish adding VIPs.
10. Alternatively, type **n** to add more VIPs and repeat Step 1 — Step 9.

Enabling the ALUA Mode

Enabling ALUA Mode results in making a SCSI target available from both nodes, even though it is physically present on only one node. The target with the node is considered active. The other node is considered standby.

When you use ALUA mode in conjunction with client side multi-pathing, it ensures target re-scanning is not required on failover, and also that the path to the standby node is valid even when not in use, (because the multi-path client continuously checks the state of the standby node and the path to the standby target).

❖ *To enable the ALUA mode, using NMV:*

1. Click **Settings > HA Cluster**.
2. Select **Advanced** > **Miscellaneous Options**.
3. Select the **Enable ALUA mode** checkbox.

   To diable the ALUA mode, remove the selection from the checkbox.

7

System Operations

This section includes the following topics:

- About System Operations
- Viewing the HA Cluster Status
- Manually Triggering a Failover
- Verifying Shared Volume Status
- Viewing Support Logs

About System Operations

There are a variety of commands and GUI screens to help you with daily cluster operations. There is a set of cluster-specific commands to supplement NMC.

Viewing the HA Cluster Status

You can view the status of the HA Cluster and heartbeats at any time.

❖ *To view the HA Cluster configuration, using NMV:*

1. In the **Cluster Settings** panel, click **Status**.
2. Click the tabs to view **Cluster Status** and **Heartbeat Status**.

❖ *To view the HA Cluster configuration, using NMC:*

1. Type:

   ```
   nmc:/$ show group rsf-cluster <cluster_group_name>
   ```

   Example:

   ```
   nmc:/$ show group rsf-cluster HA-Cluster
   ```

   System response:

   ```
   PROPERTY                      VALUE
   name                        : HA-Cluster
   appliances                  : [NodeA NodeB]
   ```

```
machinesigs             :
{"NodeA":"XXXXXXXXXX","NodeB":"YYYYYYYYYY"}
hbipifs                 : NodeA:NodeB: NodeB:NodeA:
netmon                  : 1
info                    : Nexenta HA-Cluster
generation              : 1
refresh_timestamp       : 1375745271.30001
type                    : rsf-cluster
creation                : Jan 8 22:34:50 2014

SHARED VOLUME: ha-vol
svc-ha-vol-shared-vol-name : ha-vol
svc-ha-vol-ipdevs          : ha-vol NodeA:e1000g0
NodeB:e1000g0
svc-ha-vol-ipdevs-IPv6     :
svc-ha-vol-attached-vols   :
svc-ha-vol-main-node       : NodeA
svc-ha-vol-inittimeout     : 20
svc-ha-vol-runtimeout      : 8
svc-ha-vol-mhdc-disable    : n
svc-ha-vol-monitor         :
{"NodeA":{"monitor":"","ipdevs":{"e1000g0":""}},"NodeB":{
"monitor":"","ipdevs":{"e1000g0":""}}}
svc-ha-vol-resdisks        : NodeA:c1t3d0 NodeB:c1t1d0
```

```
HA CLUSTER STATUS: HA-Cluster
NodeA:
 ha-vol       running        auto    unblocked  ha-vol       e1000g0   20  8
NodeB:
 ha-vol       stopped        auto    unblocked  ha-vol       e1000g0   20  8
```

Manually Triggering a Failover

You can manually trigger a failover between systems when needed. Performing a failover from the current appliance to the specified appliance causes the volume sharing service to stop on the current appliance, and the opposite actions take place on the passive appliance. Additionally, the volume exports to the other node.

Note: You must first set all cluster operations to manual mode.

❖ *To manually trigger a failover, using NMC:*

1. Verify that shared volume is in healthy state by typing:

   ```
   nmc:/$ zpool status <shared-volume>
   ```

   Example:

   ```
   pool: <shared-volume>
    state: ONLINE
     scan: none requested
   ```

```
config:

        NAME            STATE       READ  WRITE  CKSUM
    <shared-volume>  ONLINE         0      0      0
          mirror-0   ONLINE         0      0      0
            c1t8d0   ONLINE         0      0      0
            c1t5d0   ONLINE         0      0      0
          mirror-1   ONLINE         0      0      0
            c1t9d0   ONLINE         0      0      0
            c1t6d0   ONLINE         0      0      0

errors: No known data errors
```

Warning: If any disk drive from the shared volume is in state DEGRADED, you must replace the faulted drive(s) before executing failover. Otherwise, failover may take long time or your system may freeze.

2. Type:

   ```
   nmc:/$ setup group rsf-cluster <cluster_name> failover
   ```

Verifying Shared Volume Status

Verify the status on the shared volume service using NMV or NMC.

❖ *To view the status of a shared volume, using NMV:*

1. In the **Cluster Settings** panel, click **Status**.

❖ *To view the status of a shared volume, using NMC:*

◆ Type:

```
nmc:/$ show group rsf-cluster
```

System response:

```
HA CLUSTER STATUS: HA-Cluster
nodeA:
 vol1-114      stopped         manual   unblocked  10.3.60.134  e1000g0    20   8
nodeB:
 vol1-114      running         auto     unblocked  10.3.60.134  e1000g0    20   8
```

Viewing Support Logs

Gather the information about HA Cluster event or errors from the HA Cluster log file.

❖ *To view support logs, using NMV:*

◆ Click **View Log**.

❖ *To view support logs, using NMC:*

◆ Type:

```
nmc:/$ show group rsf-cluster <cluster_group_name> log
```

This page intentionally left blank

8

Testing and Troubleshooting

This section includes the following topics:

- Repairing a Broken Cluster Service
- Replacing a Faulted Node

Repairing a Broken Cluster Service

NexentaStor tracks various appliance components, and their state. If and when failover occurs (or any service changes to a broken state), NexentaStor sends an email to the administrator describing the event.

You can execute the repair command for a cluster service. The repair command forces the import operation of a shared volume. Therefore, the shared volume must be exported on both nodes.

Note: During the NexentaStor installation, you set up SMTP configuration and test so that you can receive emails from the appliance.

There are two broken states:

- **Broken_Safe**

  A problem occurred while starting the service on the server, but it was stopped safely and you can run it elsewhere.

- **Broken_Unsafe**

  A fatal problem occurred while starting or stopping the service on the server. The service cannot run on any other server in the cluster until it is repaired.

Warning: Manually verify and troubleshoot the volume before marking the state as repaired. Failure to do so could result in cross-mounting of the volume and lead to data corruption.

❖ *To repair a shared volume which is in broken state, using NMC:*

1. Verify that volume is exported on both HA Cluster nodes by typing:

```
nmc:/$ zpool status
```

The output should not include the information about the shared volume.

2. Repeat Step 1 on other node.
3. Execute the volume repair operation:

   ```
   nmc:/$ setup group rsf-cluster shared-volume repair
   <cluster_name> <volume_name>
   ```

   This initiates and runs the repair process.

❖ *To mark a service as repaired, using NMV:*

1. Click **Settings > HA Cluster**.
2. In the Action column, set the action to repaired.
3. Click **Confirm**.

Replacing a Faulted Node

NexentaStor provides a capability to restore a failed node in an HA Cluster, in case the state changes to `out of service`. There is no need to delete the cluster group on another node and reconfigure it and all of the cluster services.

❖ *To replace a faulted node, using NMC:*

1. Set up the same configuration on a new hardware. Verify that the following components are identical on the old and new NexentaStor nodes:
   - NexentaStor versions

     You must use identical NexentaStor versions for both HA Cluster nodes.
   - Network settings
   - SSH bindings
   - Hearbeats
2. Type:

   ```
   nmc:/$ setup group rsf-cluster <group_name> replace_node
   ```

   After executing the command, the system asks you to choose which node to exclude from the cluster and which new node to use instead. NexentaStor checks host parameters of the new node and if they match the requirements of the cluster group, the old one is replaced by a new one.

9

Upgrading HAC

This section includes the following topics:

- About Upgrading HAC
- Prerequisites
- Failing Over the Nodes
- Upgrading HA Cluster

About Upgrading HAC

Upgrading the cluster includes a few phases. You upgrade the node that does not own the volume, then you manually failover the volume to the node that was just upgraded, and then you upgrade the remaining node.

Prerequisites

The Prerequisites for the upgrade are the same as for the installation on a new system.

See Also:

- Prerequisites

Failing Over the Nodes

Before you can upgrade the HAC nodes, you must fail over the shared volume so you do not have to recreate them after the upgrade. After upgrading the node and before upgrading the other node, you failover the shared volume back to the newly upgraded node.

❖ *To fail over the active node to the passive node:*

1. Log in to the active node and type the following to fail it over to the passive node:

   ```
   nmc:/$ setup group rsf-cluster <cluster_name>
   <passive_node_name> failover
   ```

2. Set the failover mode to manual by typing:

   ```
   nmc:/$ setup group rsf-cluster <cluster_name> shared-
   volume <volume_name> manual
   ```

3. Type **Y** to confirm mode change.

Upgrading HA Cluster

When upgrading, you have an active and a passive node.

❖ *To upgrade both nodes, using NMC:*

1. Log in to the passive node and type:

   ```
   nmc:/$ setup appliance upgrade
   ```

2. After the upgrade successfully finishes, log in to the active node and type the following to fail over to the passive node:

   ```
   nmc:/$ setup group rsf-cluster <group_name>
   <shared_volume_name> <passive_node_name> failover
   ```

3. After failover finishes, the nodes swap. The active node becomes the passive node and vice versa. Type the following command on the current passive node:

   ```
   nmc:/$ setup appliance upgrade
   ```

4. Type the following to run the failover command on the current active node and thereby, make it passive again:

   ```
   nmc:/$ setup group rsf-cluster <group name> failover
   ```

Index

A

adding
- hostname 28
- hostname address 28
- shared volume 13
- virtual IP address 28

additional resources 4

B

binding
- nodes 12

C

cluster
- configuration data 19

configuration data
- cluster 19

configuring
- fibre channel
  - failover 23
- HA cluster 13
- iSCSI Targets
  - failover 21
- storage failover 19

D

defining
- netmasks 11

E

exclusive access
- to storage 2

F

failing over
- nodes 37

failover
- configuring iSCSI Targets 21
- NFS/CIFS 21
- service 3
- storage 2
- triggering
  - manually 32

failover mode
- setting 27

fibre channel
- failover configuring 23

G

guide
- audience v

H

HA cluster
- configuring 13
- upgrading 38

heartbeat interfaces
- modifying 18

heartbeat mechanism 17

hostname
- adding 28

I

implementing
- misc options 30

installing
- HA
  - new cluster 5
- HA automatically 6
- HA manually 6

M

mapping information 20

misc options
- implementing 30

modifying
- heartbeat interfaces 18

N

netmasks
- defining 11

network architecture 3
- sample 3

nexentastor
- appliances 1

NFS/CIFS failover 21

nodes
- binding 12
- failing over 37

P

prerequisites 5, 37

R

RSF-1 cluster service 2

S

setting
- failover mode 27

shared disk 17

shared volume
- adding to cluster 13
- status 33

SSH
- binding the nodes 12

storage
- failover 2
- failover configuring 19

support
- contact v

system operations 31

T

triggering
- failover 32

U

upgrading
- HAC 37, 38

V

virtual IP address
- adding 28

This page intentionally left blank

Global Headquarters
455 El Camino Real
Santa Clara, California 95050

Nexenta EMEA Headquarters
Camerastraat 8
1322 BC Almere
Netherlands

Nexenta Systems Italy
Via Vespucci 8B
26900 Lodi
Italy

Nexenta Systems China
Room 806, Hanhai Culture Building,
Chaoyang District,
Beijing, China 100020

Nexenta Systems Korea Chusik Hoesa
3001, 30F World Trade Center
511 YoungDongDa-Ro
GangNam-Gu, 135-729
Seoul, Korea

nexenta™