Hewlett Packard
Enterprise

CLI Administration Guide

Published May, 2023
Version 5.0.x

Legal Notices

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Hewlett Packard Enterprise believes in being unconditionally inclusive. If terms in this document are recognized as offensive or noninclusive, they are used only for consistency within the product. When the product is updated to remove the terms, this document will be updated.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java® and Oracle® are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

All third-party marks are property of their respective owners.

Publication Date

Tuesday May 30, 2023 02:48:31

Document ID

nnc1566944198129

Support

All documentation and knowledge base articles are available on HPE InfoSight at https://infosight.hpe.com. To register for HPE InfoSight, click the *Create Account* link on the main page.

Email: support@nimblestorage.com

For all other general support contact information, go to https://www.hpe.com/us/en/services/nimble-storage.html.

Contents

The Array Command Line Interface

This document deals with procedures to manage the array and to automate common tasks using the array OS command line interface (CLI). If you want to manage your array using the GUI, refer to the *GUI Administration Guide*. Not all procedures can be performed with both the CLI and the GUI.

Accessing the CLI

Procedure

1. Open an SSH client, such as PuTTY or OpenSSH.
2. Enter the host name or IP address of the array.
   If you are asked to accept the authorization key, type **yes**.
3. Log into the array with the username and the password you created when configuring the array.

   The connection is made and you can now run the array CLI commands. For example, typing **array --list** displays the name, serial, model, version, and status of the array:

```
$ array --list
c20-array2    AA-102081   CS220   3.1.988.0-354313-opt   reachable
```

   Typing **array --info c20-array2** displays other information about the array:

```
$ array --info c20-array2
Model: CS220
Extended Model: CS220-4G-12T-320F
Serial: AA-102081
Version: 3.1.988.0-354313-opt
All-Flash: No
Array name: c20-array2
Supported configuration: Yes
Link-local IP address: 169.254.54.168
Group ID: 3815276604473015474
Member GID: 1
Group Management IP: 10.18.120.51
1G/SFP NIC: 4/0
Total array capacity (MiB): 7606708
Total array usage (MiB): 0
Total array cache capacity (MiB): 305276
Volume usage (MiB): 0
Volume compression: 1.00X
Volume space saved (MiB): 0
Snapshot usage (MiB): 0
Snapshot compression: 1.00X
Snapshot space reduction: 1.00X
Snapshot space saved (MiB): 0
Pending Deletes (MiB): 0
Available space (MiB): 7606708
Member of pool: default
Status: reachable
```

   For more information about the command line interface, type **man** and the command of interest. For example, type **man array** for information about the array command. Also see the *Command Reference*.

CLI Commands

To see the complete list of commands for managing the Nimble Array, refer to the *Command Line Reference* and the man pages. In the Nimble documentation, bracketed subparameters are optional and unbracketed parameters are required.

You can also type ? at the command prompt to display a list of available commands, as shown in the following example:

```
Nimble OS $ ?
?                   folder              pool                timezone
alert               group               prottmpl            useradmin
array               halt                reboot              userauth
auditlog            help                route               usersession
cert                initiatorgrp        setup               vcenter
chapuser            ip                  shelf               version
ctrlr               migration           snap                vm
date                netconfig           snapcoll            vmwplugin
disk                nic                 software            vol
encryptkey          partner             sshkey              volcoll
failover            pe                  stats
fc                  perfpolicy          subnet
```

Type man or help for a command to see details for that command. For example, typing man snap shows you all the suboptions and variables for the snap command. Typing snap --help displays a less detailed list of all options.

Array Overview

The HPE Nimble Storage array seamlessly merges high-performance, compressed storage with capacity-optimized snapshot storage and WAN-efficient replication while it also serves the following functions:

- Provides a seamless combination of storage and backup with efficient disaster recovery
- Enables data restoration based on snapshots
- Offers data protection through Recovery Point Objective (RPO) and Recovery Time Objective (RTO) metrics
- Simplifies storage and snapshot management
- Binds multiple arrays into a single management group

The HPE Nimble Storage Array

HPE Nimble Storage arrays are engineered for high performance using flash, for low cost using dense, capacity-optimized disks, and for easy installation and administration.

HPE Nimble Storage solutions are built on the patented Cache Accelerated Sequential Layout (CASL™) architecture. CASL leverages the lightning-fast random read performance of flash and the cost-effective capacity of hard disk drives. Data written to the array is compressed and then stored in the disk-drive layer. HPE Nimble Storage arrays take advantage of multi-core processors to provide high-speed inline variable-block compression without introducing noticeable latency. CASL also incorporates efficiency features like cloning and integrated snapshots to store and serve more data in less space.

Data that is frequently accessed is tracked in the Nimble index, which ensures that frequently and recently accessed data is also held in the large adaptive flash layer. A copy of all data in the flash layer also remains safely in the disk-drive layer to ensure reliability, but now it can be accessed with the high performance and low latency made possible by flash technology.

Nimble Storage arrays provide:

- Higher storage efficiency to reduce the storage footprint by 30 to 75 percent
- Non-disruptive scaling to fit changing application needs through increased performance, or capacity, or both
- Maximized data and storage availability with integrated data protection and disaster recovery
- Simplified storage management and reduced day-to-day operational overhead

Nimble arrays are easy to install and manage. Consolidation of storage and backup, automatic failover, reusable schedules based on application usage, and application integration are combined with a clean, intuitive GUI that improves the ease of administration. For automation ease, a command-line interface (CLI) and RESTful API are also provided.

Array Features

Arrays provide features to enhance performance, value, and ease of use. The all-inclusive features described in this documentation do not require extra licensing.

Feature	Function	Benefit
Core Functionality		
Dynamic Caching	Reads active data from flash cache, which is populated on writes or first read	Accelerates read operations, with sub-millisecond latency
Write-Optimized Data Layout	Coalesces random writes and sequentially writes them to disk as a full stripe	Accelerates writes as much as 100x, and gets sub-millisecond latency and optimal disk utilization

Feature	Function	Benefit
Universal Compression	Always-on inline compression for all workloads	Reduces capacity needs by 30-75%, depending on the workload, with no performance impact
Thin Provisioning	Allocates disk space to a volume only as data is written; thin provision "stun" is supported to add greater flexibility	Pools storage, shares free space, and maximizes utilization
Offloaded Data Transfer (ODX) for Windows	Interacts with storage devices to move data through high-speed storage networks	Improves file copy speed
Scale Performance	Non-disruptively upgrade controllers or swap in higher capacity SSDs	Scales performance to manage large amounts of active data
Scale Capacity	Non-disruptively add external disk shelves	Increases storage capacity to 100s of TB per system
Instant Snapshot and Recovery	Backs up and restores data using point-in-time, space-efficient snapshots taken at regular intervals	Retains months of frequent snapshots, which improves RPO with no performance impact. Eliminates backup windows and speeds up restores, which improves RTO.
WAN-efficient Replication	Replicates the compressed data changes to the secondary site for disaster recovery	Deploys affordable and verifiable disaster recovery and efficiently backs up remote sites over the WAN
VLAN Tagging	Part of the 802.1Q frame header containing a VLAN ID between 1 and 4094 that identifies the VLAN to which the frame belongs	Allows a switch to forward that frame only to the appropriate VLAN
Zero-Copy Clones	Creates copies of existing active volumes without needing to copy data	Creates clones in seconds and saves disk space, which is ideal for VDI and test/development environments
Host and Application Integration		
Custom Application Profiles	Predefined policies for block size, caching, compression, and data protection for Microsoft applications and VMware virtual machines	Eliminates the need to manually tune storage and data protection configurations
Windows VSS Enablement	HPE Storage Driver for the Microsoft VSS framework for consistent backup	Takes application-consistent backups and simplifies data protection for Exchange and SQL Server
VMware Integration	Monitors, provisions, and takes snapshots from VMware vCenter	Manages storage from vCenter and takes consistent backups of virtual machines
VMware Site Recovery Manager Adapter	Supports disaster recovery automation for VMware including failover/failback	Simplifies disaster recovery, including testing failover/failback
Management and Support		
Proactive Wellness and DNA	Real-time monitoring and analysis; sends alerts on critical issues	Spots and remedies potential issues to maximize uptime, performance, and utilization (no user data is accessed or collected by DNA)

Feature	Function	Benefit
Secure Remote Support	Allows remote troubleshooting, configuration, and problem resolution	Reduces burden on IT staff and quickly resolves problems
Non-Disruptive Upgrades	Upgrades software with no disruption to applications	Maximizes uptime and user productivity through continuous availability

Note: Recovery Point Objective (RPO) and Recovery Time Objective (RTO) are among the most important parameters of a disaster recovery plan.

- RPO: The point-in-time to which systems and data must be recovered after an outage. The interval reflects the amount of data loss a business can survive.
- RTO: The amount of time it takes to recover systems and data after an outage. The interval reflects the amount of downtime a business can survive.

For RPO and RTO, a shorter time interval is better.

IOPS and MiB/s Limits on Volumes and Folders

You can set both IOPS (input/output requests per second) and MiB/s (mebibytes per second) limits when you set up a Quality of Service policy for a volume or Storage Policy-Based Management (SPBM) for Virtual Volumes.

You specify the IOPS and MBps limits separately. The input/output requests are throttled when either the IOPS limit or the MBps limit is met.

The default upper bound value for both the IOPS and the MiB/s is unlimited. The lower bound for IOPS is 256 requests. The lower bound for the MiB/s value is calculated as greater than or equal to the IOPS limit multiplied by the volume block size. This way the MiB/s value does not throttle the IOPS for the volume.

If the volume contains folders, then the IOPS and MiB/s limits work as an aggregate. The input/output requests to the volumes under the folder are throttled when the cumulative IOPS of all the volumes under that folder exceeds the folder IOPS limit or when the cumulative throughput of all the volumes under that folder exceeds the folder MiB/s limit. When this happens, all the volumes are throttled equally.

You can set these limits when you create or edit either a volume or a folder. You can set separate limits for both volumes and folders. You can view the limits for volumes and folders on the Performance tab at:

Manage > **Data Storage** > **Performance**

If you are creating or editing a folder, you can set the limits at the folder level that are in addition to the volume-level limits. The default limits are at the volume level. You can view the limits on the Folder tab at:

Manage > **Data Storage** > **Folder**

You can monitor the throughput and bandwidth for the array, the volumes, and the folders by selecting:

Monitor > **Performance**

Horizontal Scaling

Scale-out, often referred to as horizontal scaling, means adding arrays to a *group*. Performance and capacity scale linearly as you add arrays to the group. Grouping multiple arrays so that they can be managed as one entity provides significant manageability benefits, because it appears as if you are managing a single large array. Scale-out simplifies load balancing and capacity management, as well as hardware and software life-cycle management.

From an organizational standpoint, scale-out establishes a group of merged systems upon which storage pools can be developed. Volumes are created within pools that can span multiple physical arrays. A volume might exist on one array or span multiple arrays in a group by virtue of how the pool is configured.

Scale-out is a *peer-to-peer* technology where each array can operate independently, but can be managed as a single pool of storage. With scale-out, you not only add more disk and flash memory, but also CPU, system memory, network links, and so on.

Scale-out requires that all arrays in a group have the same version of array software installed.

Relationship of Groups, Pools, Arrays, Folders, and Volumes

A *group* is a collection of one to four Nimble arrays that are physically connected. Logically, they represent a single storage entity to aggregate performance and capacity, and to simplify management. All arrays in the group must be either iSCSI or Fibre Channel. A group contains one or more disjoint pools. For groups of arrays, data can be striped across the arrays in the same pool. For most administrative tasks, a group looks and feels like a single array. You administer the group by connecting to its management IP address, hosted by one of the arrays in a group.

A single-array group is formed when you configure a Nimble array.

Figure 1: Relationships of Groups, Pools, Arrays, Folders, and Volumes in Single-array Groups

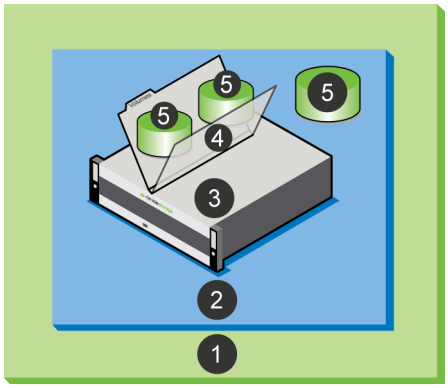

1 Group

2 Storage Pool

3 Array

4 Folder

5 Volume

A group is scaled-out, or expanded to a multi-array group, by adding unconfigured or configured arrays to the group. Adding a configured array to a group is also known as a *group merge*.

Figure 2: Relationships of Groups, Pools, Arrays, Folders, and Volumes in Multi-array Groups

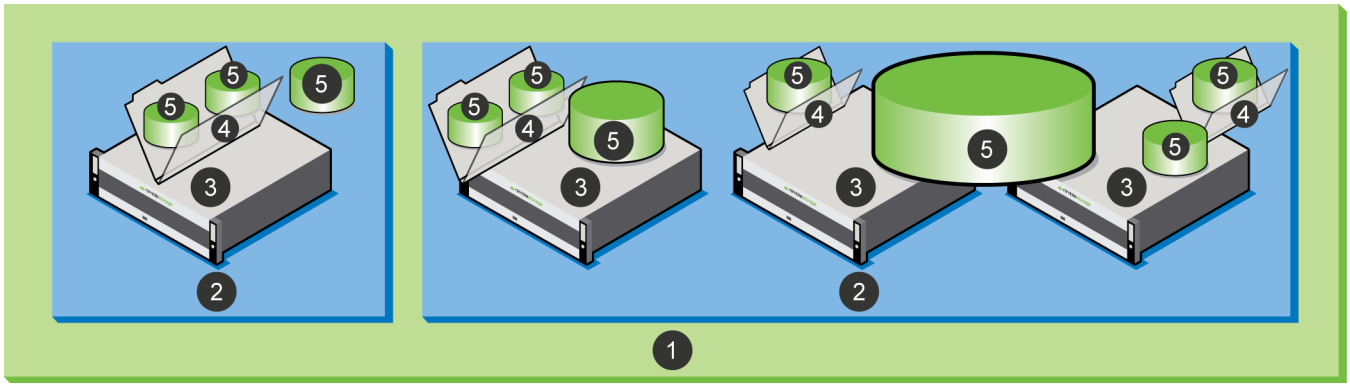

1 Group

2 Storage Pool

3 Array

4 Folder

5 Volume

You can remove an array that is part of a storage pool by evacuating the data from the array and then removing the array from the pool. When you need to remove an array from a group, you can evacuate the data to another array in the group (for example, moving a volume), provided there is enough space for the data that is stored on the array being removed.

Pools

A storage pool confines data to a subset of the arrays in a group. A storage pool is a logical container that contains one or more member arrays in which volumes reside. The system stripes and automatically rebalances data of resident volumes over the members across the pool. Storage pools dictate physical locality and striping characteristics. A member array can only be a part of one storage pool. Note the following:

- Volumes and their respective snapshots and clones reside within a pool and are tied to a specific pool.

  Volumes are also referred to as logical units (LUNs). On an array, LUNs are exposed as volumes.

- You can migrate (move) volumes between different pools.
- Volume collections are not tied to pools and can contain volumes that reside in different pools.

Use single-array pools under the following conditions:

- Fault isolation is a priority
- Linux hosts, or any hosts that do not have a supported HPE Storage Connection Manager available, can access the arrays

Consider using multi-array pools under the following conditions:

- Scaling performance and capacity, as well as consolidating management, are priorities
- Windows and ESX hosts that have a supported Connection Manager available can access the arrays

> **Note:** If your host system is not running an MPIO module, you experience a decrease in I/O performance when connected to a volume that spans multiple arrays. The drop is caused by the data paths among the arrays being redirected. Install the Connection Manager, which sets up the optimum number of iSCSI sessions (only iSCSI) and finds the best data connection to use under MPIO.

Feature Support on iSCSI and Fibre Channel Arrays

Fibre Channel arrays have slight differences in feature compatibility compared to iSCSI arrays.

The following table lists specific features and identifies whether each feature is supported on an iSCSI array and on a Fibre Channel array. For additional details, see the documentation specific to each feature.

Feature	Supported in an iSCSI Array	Supported in a Fibre Channel Array
Multi-array pools	Yes	Yes
Multi-array groups	Yes	Yes
Pool merge	Yes	Yes
Volume move	Yes	Yes
Array add	Yes	Yes
Discovery IP address	Yes	No
CHAP users	Yes	No

> **Note:** You cannot convert an iSCSI array to a Fibre Channel array or a Fibre Channel array to an iSCSI array.

Integration

For optimal integration, you should familiarize yourself with the Windows and VMware environments. Both Microsoft and VMware provide conceptual and practical information in the form of knowledge base articles, online manuals, and printed books.

About HPE Nimble Storage Windows Integration

The HPE Nimble Storage Windows Toolkit contains the necessary components to use Microsoft Volume Shadow Copy Service (VSS) to provide the backup infrastructure for Microsoft Exchange and SQL servers. It also enables you to initialize and configure a Nimble array from your Windows host.

Table 1: HPE Nimble Storage Windows Toolkit Components

Component	Description
Nimble Protection Manager	Provides the Nimble VSS requester and the Nimble VSS hardware provider. These features enable you to use the Microsoft VSS to take application-consistent snapshots on an HPE Nimble Storage array.
Nimble Connection Manager	Sets up the optimum number of iSCSI sessions and finds the best data connection to use under MPIO.
Nimble Setup Manager	Identifies uninitialized HPE Nimble Storage arrays and initializes them. Then it takes you to the NimbleOS GUI to finalize the configuration.

About HPE Nimble Storage VMware Integration

Many HPE Nimble Storage VMware integration features are preinstalled in the NimbleOS. There are some features, such as Storage Replication Adapter (SRA) and Nimble Connection Manager (NCM) for VMware, that you must install.

Table 2: HPE Nimble Storage VMware Integration Components

Component	Description
Nimble Connection Manager for VMware	Manages connections from the host to volumes on Nimble systems. NCM optimizes the number of iSCSI sessions and finds the best data connection to use under MPIO. It includes a Nimble DSM that claims and aggregates data paths for the Nimble array volumes. You must install NCM on a host running ESXi 5.0, 5.1, 5.5, 6.0, or 6.5.
vStorage APIs for Array Integration (VAAI)	Provides hardware acceleration by enabling the WRITE SAME, UNMAP, ATS,Thin Provisioning Stun, and XCOPY features.
vCenter Plugin	Allows you to create and manage datastores on the Nimble array as well as use the vCenter Server to create vCenter roles and edit protection schedules. You must register the Nimble vCenter Plugin with the vCenter Server
VMware Synchronized Snapshots	Enables application-consistent snapshots within VMware environments.
Storage Replication Adapter	Allows a Nimble array to support VMware Site Recovery (SRM) to perform array-based disaster recovery. You must install SRA on the Windows server that runs SRM.
VMware Virtual Volumes (VVols)	Enables you to use VMware virtual disks mapped to Nimble volumes on the HPE Nimble Storage array without requiring that you know the implementation details of the underlying storage. You must have vSphere 6.0 or later, ESXi 6.0 or later, and VASA Provider to use VVols. NimbleOS provides VASA Provider as part of the vCenter Server plugin.

Wellness and DNA

Diagnostics for Nimble Analytics (DNA) is the implementation of Nimble's proactive wellness feature that increases storage uptime and keeps arrays running at peak performance and efficiency. Support tools and staff monitor *heartbeats* and logs from your system and analyze a variety of parameters in real time. Any anomalies such as configuration errors or abnormal operating conditions are logged and you are alerted before a failure occurs.

These tools automatically resolve over 75 percent of issues and therefore decrease escalations. For the customer, that means faster resolution of issues. If needed, our experienced support staff can also perform secure remote troubleshooting, configuration, and problem resolution, helping resolve issues while maintaining data security. Nimble Storage non-disruptive upgrades mean you can upgrade your array and all new features and software releases with no planned downtime.

You should enable DNA when you first configure the array.

Access Controls

On your Windows server or ESXi host, you must configure iSCSI or Fibre Channel connections with your volumes on the array in order for your server or host to access those volumes. By using certain features of iSCSI or Fibre Channel, the arrray OS can control who has access to your volumes. *Access control list (ACL)* is another term for access control.

The array OS supports multiple methods of access control: user permissions, CHAP accounts, iSCSI initiator groups, Fibre Channel initiator groups, and VLAN segmentation and tagging. You can apply access controls when you create a volume or at any later time.

User Permissions

Each individual accesses the array through a user account, which is created by an Administrator. Users log into the array with their user name and password. The Administrator has two methods of access control for each user account:

- Role or permission level
- Enabling or disabling the user account

CHAP Accounts

Challenge-Handshake Authentication Protocol (CHAP) is an authentication method that servers use to verify the identity of remote clients. CHAP verifies the identity of the client by using a *handshake* when establishing the initial link and at any time afterwards. The handshake is based on the exchange of a random number known to both the client and server.

CHAP accounts are not required to establish a functional data connection between an array and a Windows server or ESXi host.

> **Note:** CHAP works only with iSCSI; it is not supported on Fibre Channel.

iSCSI Initiator Groups

An iSCSI initiator group is a collection of one or more iSCSI initiators, with each initiator having a unique iSCSI Qualified Name (IQN) and IP address. Each IQN represents a single Network Interface Card (NIC) port on an iSCSI-based client in the form of a Windows server, ESXi or Linux host. Configure iSCSI initiator groups on the array; configure client-side iSCSI initiators according to the vendor's recommendations.

> **Note:** By default, iSCSI volumes deny access to initiators. To allow initiators in an initiator group to access a volume, you must configure an ACL that includes the desired initiators and attach it to the volume.

iSCSI Initiators

The array uses iSCSI initiators and iSCSI targets as a method of communication. The volumes and snapshots on the array act as targets, or data providers. This provides an extra level of security.

A common method for limiting access is to create groups of iSCSI initiators that are managed as a single unit. The iSCSI groups can be based on OS, application, or any other logical common requirement. All members of the iSCSI group are granted access to the target volume.

To access an iSCSI target (the volumes and snapshots) you must have a supported iSCSI initiator installed on a computer within your network subsystem. You then grant the initiator permission to access the array's iSCSI discovery IP address.

iSCSI Targets

To iSCSI initiators, the volumes appear as iSCSI targets, allowing them to send and receive traffic to the volumes.

When you use iSCSI initiator discovery, all volumes appear in the list of discovered targets. You will notice another entity, named **control**. The control entity is used by the array to facilitate using a shared IP address for port pairs. Do not connect to the control entity.

Typically, multiple clients (iSCSI initiators) do not connect to the same volume. Except in cases where multiple client access is expected (such as with clustered applications) doing so could cause data corruption if both clients attempt to write data simultaneously.

If an iSCSI target appears to be stuck in a reconnecting state, remove the target and refresh the list,then reconnect the target.

> **Note:** (Windows Only) Volumes of more than 2 TB must be formatted using the GPT (GUID Partition Table). Do not format them as MBR (Master Boot Record). If you are using a Windows Vista client, only 2TB will be imported and the remaining capacity will be lost. This is a limit of the MBR partition. For a detailed discussion of GPT, see the GUID Partition Table entry in Wikipedia, or the Windows and GPT FAQ at http://www.microsoft.com> Device Fundamentals > Storage.

Fibre Channel Initiator Groups

A Fibre Channel initiator group is a collection of one or more initiators, with each initiator having a unique World Wide Port Name (WWPN). Each WWPN represents a single Host Bus Adapter (HBA) port on a Fibre Channel-based client in the form of a Windows server, ESXi or Linux host. Configure Fibre Channel initiator groups on the Nimble array; configure client-side Fibre Channel initiators according to the vendor's recommendations.

> **Note:** By default, Fibre Channel volumes deny access to initiators. To allow initiators in an initiator group to access a volume, you must configure an ACL that includes the desired initiators and attach it to the volume.

VLAN Segmentation and Tagging

VLANs provide logical segmentation of networks by creating separate broadcast domains, which can span multiple physical network segments. Arrays support VLANs, after the initial array setup has been completed. VLANs can be grouped by departments, such as *engineering* and *accounting*, or by projects, such as *release1* and *release2*. Because physical proximity of the end-stations is not essential in a VLAN, you can disperse the end-stations geographically and still contain the broadcast domain in a switched network. You can also use a Target Subnet List to limit the number of VLAN-tagged subnets that can access an end-station. VLANs provide a number of advantages such as ease of administration, confinement of broadcast domains, reduced network traffic, and enforcement of security policies.

VLAN tagging simplifies network management by allowing multiple broadcast domains or VLANs to be connected through a single cable. This is done by prepending a header to each network frame, which identifies the VLAN to which the frame belongs. Switches can be configured to route traffic for a given VLAN through a certain set of ports that match the VLAN tag.

VLANs and Initiator Groups

An initiator group is a group of one or more initiator names (iSCSI IQNs or Fibre Channel WWPNs) that can be used to grant access to volumes or LUNs in a SAN, or to assign those volumes or LUNs to a VLAN.

The array OS supports the configuration of multiple initiator groups with access control. Access can be added to or removed from a volume by configuring initiator groups. By default, iSCSI volumes are set up with full access, while Fibre Channel volumes are set up with no access.

The array OS also supports the configuration of an initiator group with no access control, by entering "*" in both the IQN and IP fields in the Edit an Initiator Group screen. This configuration allows the initiator group to control access through the Target Subnet List, if selected.

VLANs allow you to have multiple subnets per interface. By selecting the Target Subnet List, you can limit the number of subnets that have access to a volume. This is useful when there are so many subnets that timeouts may occur, for example when subnets need to be scanned on volume restart. It can also be used for security, to prevent certain subnets from accessing the volume.

Major Workflows

The following major workflows tables describe the steps necessary to complete commonly performed end-to-end, multi-task configurations. Each workflow table provides:

- A hyperlinked list of steps (tasks) required to complete each workflow
- Descriptions and guidelines for each step in the workflow

Updating array OS Workflow

The array OS is the software that runs on the array. HPE Storage provides regular maintenance releases and periodic updates to the array OS. Before you configure your array, be sure you have the latest version of the array OS installed.

If an array OS update is available for your array, you can download it now and install it later. The CLI has an automated procedure to install updates.

The following table describes the workflow for updating the array OS software on an array:

Step	Notes
1. Find the Array OS Version on page 25	You must note the version of the array OS currently running on your array to help you select the right update.
2. Update the Array OS on page 26	Installs the array OS software update on the array. To perform this task, the array must have access to the HPE InfoSight over an Internet connection. **Note:** If your array does not have access to HPE InfoSight, you can still manually download the array OS software from a PC or other workstation with an Internet connection, then update the array OS using the GUI. For more information, refer to the *GUI Administration Guide*.
3. Verify the Array OS Update on page 28	After the aray OS update has finished, verify that the update was successful.

Provisioning Storage Volumes and Performance Policies Workflow

The following table describes the workflow for provisioning volumes and performance policies on your array:

Step	Notes
1. Create a Volume on page 64	You can create either iSCSI or Fibre Channel volumes. When a volume is created, the only mandatory options are the name and the volume size. Additional settings, for example, the threshold values for volume and snap-shot reserves, quotas, and warning levels, can be set after the volume is created using the --edit option with the **vol** command.
2. Create a Performance Policy on page 72	A performance policy helps optimize the performance of the volume based on the characteristics of the application using the volume.

Setting Up Replication Workflow

The following table describes the workflow for setting up replication on your volume collections:

Step		Notes
1		Create a replication partner and add it to the volume collection.
2	Configure Bandwidth Limitations for a Replication Partner on page 114	
3	Test the Connection between Replication Partners on page 107	
4	Add Replication to a Volume Collection on page 112	(Optional) Add replication to the volume collections that are replication partners, if replication is not already enabled. If replication is already enabled on both volume collections, you can skip this step.
5	Perform a Volume Collection Handover on page 113	

Creating a vVol Datastore Workflow

VMware uses virtual volumes (vVols) to manage virtual machines (VMs) and their data (such as VMDKs and physical disks). Supporting vVols enables a volume to reside in a vVol datastore that maps to a folder on an array. A folder can contain both vVols and regular volumes.

vVols are visible in the CLI as regular volumes, where you can monitor their capacity and performance. However, you must use the vCenter UI to manage vVols.

> **Note:** Before configuring vVols, you must have vCenter Server installed and the HPE vCenter Plugin registered with the server. For more information, refer to the VMware vCenter installation documentation.

The following table describes the workflow for creating a vVol datastore.

Keep the following in mind as you perform the steps in the table.

- The database and log files need to be on a separate VMDK.
- You cannot use the VMware synchronous replication feature for volume collections.
- The VSS option to quiecse the operating system is a Microsoft VSS function, not an HPE Storage VSS function.

Step		Interface	Notes
1	Register a vCenter Plugin with vCenter Server on page 75	Array OS	Specify the **--extension vasa** command option to register the vCenter extension for a VASA provider.
2	Create a Folder on page 78	Array OS	Specify the **--agent_type vvol** command option to create a folder for a vVol datastore.
3	Create a vVol Datastore	vCenter	For more information about creating a vVol datastore in a folder, refer to the *VMware Integration Guide*.

Restoring Snapshot Data from Clones Workflow

In the rare case that an entire dataset is corrupted, you can restore the entire volume. Restoring a volume from a snapshot replaces the data in the volume with the data that existed when you created the snapshot. Restoring a volume does not destroy the existing snapshot.

The following table describes the workflow for restoring snapshot data from a clone:

Step		Notes
1.	Change a Volume's State on page 67	Before restoring data from a snapshot, the volume to be restored needs to be taken offline.
2.	Take a Manual Snapshot on page 99	Taking a manual snapshot of the volume saves its most recent state.
3.	Clone a Snapshot on page 99	Clone the snapshot volume from which to restore the data.
4.	Restore a Volume from a Snapshot on page 68	Restoring the volume from the snapshot restores the snapshot data to the volume.
5.	Change a Volume's State on page 67	After the volume is restored from the clone, you need to set the volume back online.

Hardware and Software Updates

There are several ways to keep an array up-to-date, to improve its performance, and to increase data storage capacity. (The term *upgrade* refers to array hardware.) The term *update* refers to the software.

Upgrades and Updates

There are multiple upgrade paths for hardware, as well as software updates available that you can perform on the array, depending on your array model. If you want to increase data storage capacity, you can add expansion shelves without having to perform an update or upgrade.

Hardware Upgrades

Depending on which model of storage array you have, there are multiple upgrade paths available. Cache, controllers, PCI devices, and capacity can be upgraded independently from each other.

Details about possible upgrades can be found in the compatibility matrix, which can be accessed on InfoSight.

Contact your sales rep when you want to upgrade.

See the applicable upgrade quick start guide that ships with the upgrade component. The Hardware Guide for your array model also covers upgrades.

Updates

The NimbleOS is the software that runs on the array. HPE Nimble Storage provides regular maintenance releases and periodic updates to the NimbleOS. Maintenance releases typically correct bugs and enhance features. Updates involve a major new release of the NimbleOS with new features and capabilities.

If a NimbleOS update is available for your array, you can accomplish the update in less then an hour for each array in the group. (If you want to combine your HPE Nimble Storage arrays as members of a group, each of them must have the same version of NimbleOS 2.x or later installed.) The update procedure works on one controller at a time and results in a controller failover. To avoid any data service disruption, make sure that the initiators connected to the array have proper MPIO timeouts configured before performing the software update.

If the HPE Nimble Storage Windows Toolkit (NWT) is not installed on the Windows hosts, be sure to configure timeout values appropriately. See Timeout Values on page 226.

The NimbleOS has an automated procedure to download and install updates. Or, you can download NimbleOS software at the HPE InfoSight™ at https://infosight.hpe.com. If you do not have a user account, you can create one on your first visit.

Find the Array OS Version

You must note the array OS version currently running on your array to help you select the right update. Use the CLI to locate the array OS software version.

Procedure

1. (Optional) Identify the name of the array.

   array --list

2. Display the array information.

   array --info *name*

3. Note the array OS version currently running on your array.

Example

Identifying an array and listing its information to find its array OS version.

```
$ array --list
---------------------+-----------------+-----------+------------------+------
-----
Name                  Serial             Model       Version           Status
---------------------+-----------------+-----------+------------------+------
-----
rack6array1          rack6array1      vmware     0.0.0.0-199310-opt reachable

rack6array2          rack6array2      vmware     0.0.0.0-199310-opt reachable

rack6array3          rack6array3      vmware     0.0.0.0-199310-opt reachable

$ array --info rack6array3
Model: vmware
Extended Model: vmware-4G-5T-160F
Serial: rack6array3
Version: 0.0.0.0-199310-opt
All-Flash: No
Array name: rack6array3
Supported configuration: Yes
Link-local IP address: 192.168.1.186
Group ID: 6489051080133784665
Member GID: 1
Group Management IP: 10.10.164.191
1G/10G_T/SFP/FC NIC: 4/0/0/0
Total array capacity (MiB): 23980
Total array cache capacity (MiB): 16384
Volume usage (MiB): 0
Volume compression: N/A
Volume space saved (MiB): 0
Snapshot usage (MiB): 0
Snapshot compression: N/A
Snapshot space reduction: N/A
Snapshot space saved (MiB): 0
Pending Deletes (MiB): 0
Available space (MiB): 23980
Member of pool: default
Status: reachable
```

Update the Array OS

There are regular maintenance releases and periodic updates to the array OS. Before you configure your array, be sure you have the latest array OS version installed.

Before you begin

To perform this task, the array must have access to the HPE InfoSight over an Internet connection.

> **Note:** If your array does not have access to HPE InfoSight, you can still manually download the array OS software from a PC or other workstation with an Internet connection, then update the array OS using the GUI. For more information, refer to the *GUI Administration Guide*.

Procedure

1. Use a secure shell (SSH) utility to log in to the array or group leader.

Note: You must log in with a Power User or Administrator account.

2. View a list of array OS download files:

   software --list

   A list appears with the array OS versions available to the array.

   Note the Version numbers and Status terms. The higher the number, the newer the version of the download file. See the Status terms in the following table.

   Table 3: Array OS download files

Status Term	Description	Application
available	Newer version of the array OS than the one on the array.	Updates your array.
installed	Same version as the array OS running on the array.	No change.
rollback	Older version of the array OS than the one on the array.	Rolls back to an earlier version. Normally used for troubleshooting.

   Multiple array OS download files might be marked *available*. Be sure to download the latest file.

3. Download the appropriate array OS update.

   software --download *version*

4. Verify the software download before installing the software.

   software --precheck

5. Start the software update and Accept the End User License Agreement (EULA).

   software --update --accept_license

   The update begins as soon as you accept the EULA.

   The array OS update process takes about 20 minutes per array. During that time, a controller failover and a browser reload occur automatically. The array itself remains online and available throughout the update.

   If you have multiple arrays in a storage group, all arrays in the group are updated, one at a time, to the same array OS version.

   Note: If your connection to the array drops during the update, you might not be able to reestablish until the update is done.

6. (Optional) Monitor the status of the update:

   software --update_status [**--verbose**]

   When the update is finished, the new array OS version is listed as the Current version.

Example

```
$ software --list
-----------------+-----------+-------------------------
Version           Status      Size (MiB)
-----------------+-----------+-------------------------
2.1.0.0-38453-opt available       843
2.1.0.0-29743-opt available       839
2.1.0.0-27118-opt available       837
2.1.0.0-24696-opt rollback        839
$ software --download 2.1.0.0-38453-opt
INFO: Download of software package version 2.1.0.0-38453-opt started.
```

```
Use software --download_status command to check status.
$ software --precheck
INFO: Software Update precheck passed.
$ software --update --accept_license
software --update_status
Updating group to version: 2.1.0.0-38453-opt
Group update start time Nov 19 2015 12:54:36
Group update end time N/A
Updating array: array1
Array update status: 1 of 1: Controller A is unpacking update package
```

Verify the Array OS Update

Procedure

1. List the alerts and events.

   alert --list

2. Verify the array OS update by looking for the *6003 update* events in the log. For example:

```
9275 INFO Nov 19 2016 16:21:38  6003 update  AC-102566  Successfully updated
 software to version 3.2.0.0-38453-opt on controller A
9280 INFO Nov 19 2016 16:24:20  6003 update  AC-102566  Successfully updated
 software to version 3.2.0.0-38453-opt on controller B
```

   Two 6003 update events, one for each controller, indicates a successful software update.

Network Configuration

A network configuration enables an array to be accessed and managed from the network, communicate with other arrays in a group, carry data traffic, and replicate volumes. It contains all the network parameter settings on an array, including:

- Network configuration profiles
- IP addresses
- Subnets
- Routes
- Network interfaces
- iSCSI connection
- Fibre Channel connection
- VLANs and VLAN tagging

For more information about network considerations during array installation (depending on your topology), refer to the Installation Guide or Hardware Guide for your array model.

Network Configuration Profiles

You assign network settings to one of three network configuration profiles: Active, Backup, and Draft. You can make changes to (edit) all three profiles while the array is running.

You can create a Draft profile from an Active or a Backup profile. After you have finished creating a new network configuration using the Draft profile, you can promote it to be the Active profile.

When the Active profile is revised, by being edited or replaced by the Draft configuration, the previous Active profile becomes the Backup profile.

Draft Network Configuration Profile

When making changes to the Active and Backup network configuration profiles using the CLI, each command is validated as soon as it is commited, which can sometimes cause errors if you commit commands in the wrong order, or if there is an error in the command. However, the Draft network configuration profile does not validate the commands as they are committed. Instead, the commands are simultaneously validated when the profile is validated.

Validating multiple CLI commands at the same time allows you to easily build and validate a configuration without having to worry about the order in which the commands need to be committed. And if the profile validation returns any errors in your configuration, you can fix them in the Draft profile. Then, when you activate the Draft profile, you can be assured that the configuration changes will work properly. Use the following workflow when building and validating a network configuration:

- Create a Draft profile from the Active profile
- Commit commands to the Draft profile
- Validate the Draft profile
- Activate the Draft profile

Create a Draft Network Configuration Profile

You can create a Draft network configuration profile from either the Active or Backup profile.

Procedure

Create a new Draft network configuration profile from either the Active or Backup profile.

netconfig --create_draft_from {**active** | **backup**}

Example

Creating a Draft profile from the Active profile:

```
$ netconfig --create_draft_from active
```

Creating a Draft profile from the Backup profile:

```
$ netconfig --create_draft_from backup
```

Validate a Network Configuration Profile

Before you activate changes you make to a network configuration, you can validate the changes to ensure there are no errors.

Procedure

1. Change or add any IP addresses, subnets, interfaces, or routes associated with the network configuration profile to be validated.
2. Validate the network configuration profile.

   netconfig --validate {**active** | **backup** | **draft**}

Example

Validating the Draft profile after adding two IP addresses:

```
$ ip --add 192.168.80.50 --netconfig draft --type data
$ ip --add 192.168.90.60 --netconfig draft --type data
$ netconfig --validate draft
INFO: Configuration is valid.
```

Activate a Network Configuration Profile

You can activate (commit) changes made to a network configuration so that they become part of the actively running network configuration.

Procedure

1. Validate the network configuration profile to ensure that it is valid.

   netconfig --validate {**Active** | **backup** | **draft**}
2. Activate the network configuration profile.

   netconfig --activate {**Active** | **backup** | **draft**}

Example

Activating the Draft profile after validating it:

```
$ netconfig --validate draft
INFO: Configuration is valid.
netconfig --activate draft
```

Delete a Network Configuration Profile

Note: The Active network configuration profile cannot be deleted; only the Draft and Backup profiles can.

Procedure

Delete the Draft or Backup network configuration profile.

netconfig --delete {**backup** | **draft**}

Example

Deleting the Draft profile:

```
$ netconfig --delete draft
```

IP Addresses

An IP address is a 32-bit identifier for devices on a TCP/IP network. IP addresses allow devices on a network, such as servers, switches, and arrays, to communicate with each other. HPE storage arrays use IP addresses for the following purposes:

Table 4: Types of IP addresses

IP Address	Purpose
Management	Typically defined on eth1 or on eth1 and eth2 interface, the management IP address provides access to the management interface (GUI, CLI, or API) for the array group. It is also used for volume replication. It resides on the group management subnet and floats across all management only (Mgmt only) and management + data (Mgmt + Data) interfaces.
Discovery	For iSCSI arrays, each subnet has its own discovery IP address. It enables the iSCSI initiator to discover iSCSI targets for the volumes on the array. You can use this IP address for data as well as management in a single shared network. **Note:** Discovery IP addresses are not required for Fibre Channel arrays.
Data	One or more IP addresses can be configured to carry data traffic. One data IP address can be configured for each interface pair (corresponding interfaces on the two controllers). Both controllers use the same IP address but never at the same time because only one controller is active at a time. Other data IP addresses can be configured on different subnets. **Note:** In a dedicated network topology, the data IP addresses cannot be the same as the management/iSCSI discovery IP addresses.
Support	Each controller on an array must have a dedicated support IP address, which can be used for troubleshooting and technical support purposes in the event that a controller is not reachable through the management IP address. The support IP addresses must be placed on the group management subnet.

Add an IP Address

You can add management, data, discovery, and support IP addresses using the Draft network configuration profile. The Draft profile is used to validate the configuration changes.

Note: You are not required to configure all IP addresses at the same time, nor configure them in the order shown in this task.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Add a managment IP address.

   ip --add *ip-addr* **--netconfig draft** [**--array** *name*]**--type management** [**--nic** *name*]

3. Add a data IP address.

   ip --add *ip-addr* **--netconfig draft** [**--array** *name*]**--type data** [**--nic** *name*]

4. Add a support IP address.

   ip --add *ip-addr* **--netconfig draft** [**--array** *name*]**--type support** [**--nic** *name*] [**--ctrlr** {**a** | **b**}]

   Repeat this step if you want to create a support IP address for the second controller.

5. Validate the Draft network configuration profile.

   netconfig --validate draft

   Nimble OS validates the configuration. If an error exists, Nimble OS returns an error message. Resolve all errors before proceeding to the next step.

6. Activate the Draft network configuration profile.

   netconfig --activate draft --force_ip_update

> **Note:** The **--force_ip_update** option is required to activate configurations containing discover IP addresses or data IP adresses. Without this option, any attempt to update discovery IP addresses or data IP adresses in the Active network configuration profile will fail.

Example

Configuring management, data, and support IP addresses using the Draft network configuration profile:

```
Nimble OS $ netconfig --create_draft_from active
Nimble OS $ ip --add 192.168.120.55 --netconfig draft --type management
Nimble OS $ ip --add 192.168.120.56 --netconfig draft --type data
Nimble OS $ ip --add 192.168.120.57 --netconfig draft --type support --nic eth1
 --ctrlr a
Nimble OS $ ip --add 192.168.120.58 --netconfig draft --type support --nic eth1
 --ctrlr b
Nimble OS $ netconfig --validate
INFO: Configuration is valid.
Nimble OS $ netconfig --activate --force_ip_update
```

Edit an IP address

You can edit one or more existing IP addresses using the Draft network configuration profile. The Draft profile is used to validate and activate the configuration changes.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Edit an existing IP address.

   ip --edit *ip-addr* **--netconfig draft** [**--array** *name*]**--type** {**data** | **discovery** | **management** | **support**} [**--nic** *name*] [**--newaddr** *ip-addr*] [**--ctrlr** {**A** | **B**}]

3. (Optional) Repeat Step 2 to edit additional IP addresses.
4. Validate the Draft network configuration profile.

   netconfig --validate draft

   Nimble OS validates the configuration. If an error exists, Nimble OS returns an error message. Resolve all errors before proceeding to the next step.

5. Activate the Draft network configuration profile.

   netconfig --activate draft --force_ip_update

Note: The **--force_ip_update** option is required to activate configurations containing discover IP addresses or data IP adresses. Without this option, any attempt to update discovery IP addresses or data IP adresses in the Active network configuration profile will fail.

Example

Editing a management IP addresses using the Draft network configuration profile:

```
Nimble OS $ netconfig --create_draft_from active
Nimble OS $ ip --edit 192.168.120.55 --netconfig draft --type management --
newaddr 192.168.120.201
Nimble OS $ netconfig --validate
INFO: Configuration is valid.
Nimble OS $ netconfig --activate --force_ip_update
```

Delete an IP Address

You can delete one or more existing IP addresses using the Draft network configuration profile. The Draft profile is used to validate and activate the configuration changes.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Delete an existing IP address.

   ip --delete *ip-addr* **--netconfig draft** [**--array** *name*]**--type** {**data** | **discovery** | **management** | **support**} [**--ctrlr** {**a** | **b**}]

3. (Optional) Repeat Step 2 to delete additional IP addresses.
4. Validate the Draft network configuration profile.

   netconfig --validate draft

   Nimble OS validates the configuration. If an error exists, Nimble OS returns an error message. Resolve all errors before proceeding to the next step.

5. Activate the Draft network configuration profile.

   netconfig --activate draft --force_ip_update

   Note: The **--force_ip_update** option is required to activate configurations containing discover IP addresses or data IP adresses. Without this option, any attempt to update discovery IP addresses or data IP adresses in the Active network configuration profile will fail.

Example

Deleteting a management IP addresses using the Draft network configuration profile:

```
Nimble OS $ netconfig --create_draft_from active
Nimble OS $ ip --delete 192.168.120.55 --netconfig draft --type management
Nimble OS $ netconfig --validate
INFO: Configuration is valid.
Nimble OS $ netconfig --activate --force_ip_update
```

Subnets

A subnet is logical subdivision of a network. It is defined by the first IP address in the network and a netmask that specifies a contiguous range of IP addresses within that network. A subnet can be assigned to one or more network interfaces.

The maximum Transmission Unit (MTU) can be set for a subnet so that it uses either a standard, jumbo, or custom frame size. If you choose to use a custom frame size, you must specify the size in bytes.

Specifying a VLAN ID on a subnet allows an interface to be assigned to more than one subnet using IEEE 802.1Q tagged frames. Switch port configuration must match the VLAN IDs configured on the subnets for tagged assignments. For more information, refer to the procedure to Configure VLAN Tagging in the *GUI Administration Guide* or *CLI Administration Guide*.

Note: The arrays in a group communicate with each other on the "native vlan". The native vlan can be enabled on any subnet; however, there needs to be at least one subnet in which "untagged" traffic is allowed. This native vlan is used when merging a group, adding and removing an array, updating a group configuration, and updating network configurations.

Subnet Traffic Types

Traffic types are used to segregate network traffic into different subnets. A subnet can carry one of the following traffic types.

Table 5: Traffic Types

Traffic Type	Description
Management (Mgmt only)	The subnet carries only management traffic.
Data (Data only)	The subnet carries only data traffic.
Management and Data (Mgmt + Data)	The subnet carries both management and data traffic.

Subnet Traffic Assignments

Traffic assignments determine what type of iSCSI traffic a subnet carries. A subnet can have one of the following traffic assignments.

Note: Traffic assignments are not required for Fibre Channel arrays.

Table 6: Traffic Assignments

Traffic Assignment	Description
iSCSI + Group	The subnet carries both iSCSI data traffic and intra-group communication (traffic between arrays in a group).
iSCSI only	The subnet carries only iSCSI data traffic.
Group only	The subnet carries intra-group communication traffic.

IP Address Zones in Subnets

An IP address zone is a group of host IP addresses and array data IP addresses in a subnet. When using two switches for iSCSI traffic, hosts can achieve better performance by establishing iSCSI connections with data IP addresses inside the same zone, as opposed to establishing iSCSI connections with data IP addresses in a different zone.

Note: IP address zones are not required for Fibre Channel arrays.

The IP addresses within a subnet can be divided into the following IP address zone types:

Table 7: IP Address Zones Types

Zone Type	Description
None	Used for non-iSCSI enabled subnets.
Single	All IP addresses are in one zone. This is the default zoning setting. With two network switches, iSCSI connections can be routed over the inter-switch link.
Bisect	One zone includes the IP addresses from the top half of the subnet; for example, 192.168.1.128 to 192.168.1.254. The other zone takes the IP addresses from the bottom half of the subnet; for example, 192.168.1.1 to 192.168.1.127
Even/Odd	The IP addresses are grouped by their last bit. One zone includes the even-numbered IP addresses, such as 192.168.1.2, 192.168.1.4, 192.168.1.6, and so on. The other zone includes the odd-numbered IP addresses, such as 192.168.1.1, 192.168.1.3, 192.168.1.5, and so on.

IP address zones are useful for configurations that use two switches, where you want to establish connections that avoid the Inter-Switch Link. For IP address zones to work, the host and the array must have its data IP addresses configured with half of its IP addresses from one zone connected to one switch and the other half of its IP addresses from the other zone connected to the other switch. For example, assume that:

- There is single subnet, 192.168.1.0/24.
- There are two zones, defined as Red and Blue.
- Red zone consists of:
  - Host IP 192.168.1.1
  - Array Data IP 192.168.1.3
  - Array Data IP 192.168.1.5
- Blue zone consists of:
  - Host IP 192.168.1.2
  - Array Data IP 192.168.1.4
  - Array Data IP 192.168.1.6

In the IP Address Zone, the host IP addresses in the Red zone only establish connections with the data IP addresses in the Red zone. And the host IP addresses in the Blue zone only establish connections with the data IP addresses in the Blue zone. In this way, iSCSI connections do not use inter-switch link and thereby maximize I/O performance.

Figure 3: IP Address Zones

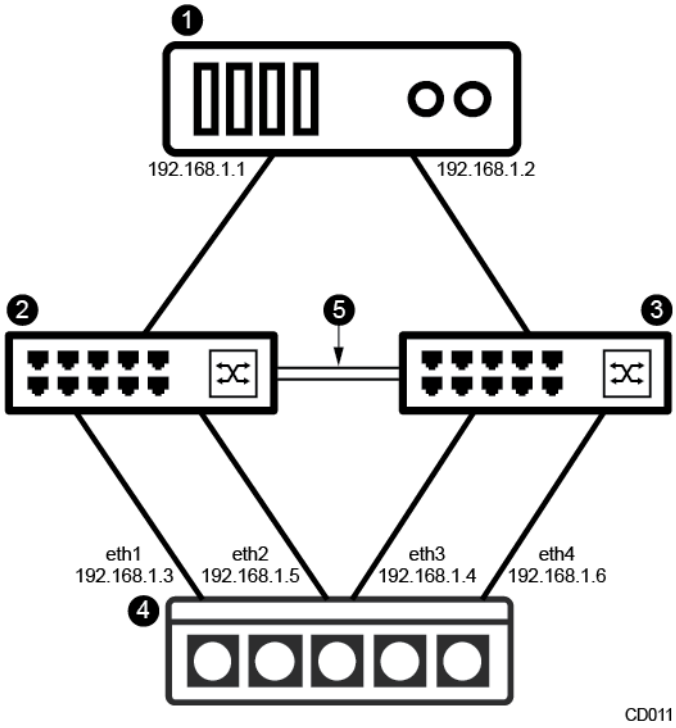

1 Host

2 Switch 1

3 Switch 2

4 Array

5 Inter-switch link

Add a Subnet

You can add one or more subnets using the Draft network configuration profile. The Draft profile is used to validate and activate the configuration changes.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Create a subnet.

   subnet --add *label* **--subnet_addr** *network_ipaddr/netmask* [**--discovery_ipaddr** *ipaddr*] [**--type** {**mgmt** | **data** | **mgmt,data**}] [**--subtype** {**iscsi**|**group**}] [**--netzone_type** {**evenodd** | **bisect** | **single**}] [**--netconfig** *name*] [**--vlanid** *id*] [**--mtu** *mtu*] [**--failover** *{enable|disable}*]

3. (Optional) Repeat Step 2 to add additional subnets.
4. Validate the Draft network configuration profile.

   netconfig --validate draft

   The array OS validates the configuration. If an error exists, The array OS returns an error message. Resolve all errors before proceeding to the next step.

5. Activate the Draft network configuration profile.

   netconfig --activate draft

Example

Configuring the data1 subnet using the Draft network configuration profile:

```
$ netconfig --create_draft_from active
$ subnet --add data1 --subnet_addr 192.168.120.0/24 --netconfig draft
$ netconfig --validate
INFO: Configuration is valid.
$ netconfig --activate
```

Edit a Subnet

You can edit one or more existing subnets using the Draft network configuration profile. The Draft profile is usesd to validate and activate the configuration changes.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Edit an existing subnet.

   subnet --edit *label* [**--new_label** *name*] [**--subnet_addr** *network_ipaddr/netmask*] [**--discovery_ipaddr** *ipaddr*] [**--type** {**mgmt** | **data** | **mgmt,data**}] [**--subtype** {**iscsi**|**group**}] [**--netzone_type** {**evenodd** | **bisect** | **single**}] [**--netconfig** *name*] [**--vlanid** *id*] [**--mtu** *mtu*] [**--failover** *{enable|disable}*]

3. (Optional) Repeat Step 2 to edit additional subnets.
4. Validate the Draft network configuration profile.

   netconfig --validate draft

   The array OS validates the configuration. If an error exists, The array OS returns an error message. Resolve all errors before proceeding to the next step.

5. Activate the Draft network configuration profile.

   netconfig --activate draft

Example

Editing the data1 subnet using the Draft network configuration profile:

```
$ netconfig --create_draft_from active
$ subnet --edit data1 --new_label data2 --subnet_addr 200.200.200.0/24 --net▶
config draft
$ netconfig --validate
INFO: Configuration is valid.
$ netconfig --activate
```

Remove a Subnet

You can remove (delete) one or more subnets using the Draft network configuration profile. The Draft profile is used to validate and activate the configuration changes.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Delete a subnet.

   subnet --remove *label* [**--netconfig** *name*] [**--force_repl**] [**--force_ip_update**] [**--force_unassign_nics**] [**--force_initiator_groups**]

3. (Optional) Repeat Step 2 to remove additional subnets.

4. Validate the Draft network configuration profile.

   netconfig --validate draft

   The array OS validates the configuration. If an error exists, the array OS returns an error message. Resolve all errors before proceeding to the next step.
5. Activate the Draft network configuration profile.

   netconfig --activate draft

Example

Removing the data1 subnet using the Draft network configuration profile:

```
$ netconfig --create_draft_from active
$ subnet --remove data1 --netconfig draft --force_ip_update
$ netconfig --validate
INFO: Configuration is valid.
$ netconfig --activate
```

Routes

Routes are paths from one network location to another; static routes are paths that do not dynamically change with changing network conditions. You can add static routes to a network configuration. For example, if you want an array to use a specific path through the network to reach a gateway, you can create a static route to the gateway.

To create a static route, you specify a subnet (network address and netmask) and the gateway IP address within the subnet.

> **Note:** The default gateway IP address is in the same subnet as the management IP address.

Add a Static Route

You can add one or more static routes using the Draft network configuration profile. The Draft profile is used to validate and activate the configuration changes.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active
2. Add a static route.

   route --add *network* **--netconfig name** [**--gateway** *gateway*]
3. (Optional) Repeat Step 2 to add more routes.
4. Validate the Draft network configuration profile.

   netconfig --validate draft

   The array OS validates the configuration. If an error exists, array OS returns an error message. Resolve all errors before proceeding to the next step.
5. Activate the Draft network configuration profile.

   netconfig --activate draft

Example

Adding a static route using the Draft network configuration profile:

```
$ netconfig --create_draft_from active
$ route --add 192.168.50.0 --netconfig draft --gateway 192.168.50.101
$ netconfig --validate
```

```
INFO: Configuration is valid.
$ netconfig --activate
```

Edit a Static Route

You can edit one or more existing static routes using the Draft network configuration profile. The Draft profile is used to validate and activate the configuration changes.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Edit an existing static route.

   route --edit *network* **--netconfig draft** [**--network** *new_network*] [**--gateway** *gateway*]

3. (Optional) Repeat Step 2 to edit more routes.
4. Validate the Draft network configuration profile.

   netconfig --validate draft

   The array OS validates the configuration. If an error exists, the array OS returns an error message. Resolve all errors before proceeding to the next step.

5. Activate the Draft network configuration profile.

   netconfig --activate draft

Example

Editing a static route network IP address and gateway IP address using the Draft network configuration profile:

```
$ netconfig --create_draft_from active
$ route --edit 192.168.50.0 --netconfig draft --network 10.190.25.0 --gateway
 10.190.25.101
$ netconfig --validate
INFO: Configuration is valid.
$ netconfig --activate
```

Delete a Static Route

You can delete one or more static routes using the Draft network configuration profile. The Draft profile is used to validate and activate the configuration changes.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Edit an existing static route.

   route --delete *network* **--netconfig draft**

3. (Optional) Repeat Step 2 to delete more routes.
4. Validate the Draft network configuration profile.

   netconfig --validate draft

   The array OS validates the configuration. If an error exists, the array OS returns an error message. Resolve all errors before proceeding to the next step.

5. Activate the Draft network configuration profile.

   netconfig --activate draft

Example

Deleting a static route using the Draft network configuration profile:

```
$ netconfig --create_draft_from active
$ route --delete 192.168.50.0 --netconfig draft
$ netconfig --validate
INFO: Configuration is valid.
$ netconfig --activate
```

Configure the Default Gateway

You can configure the default gateway by adding a route for the management subnet.

Procedure

Add a static route for the management subnet and specify the default gateway IP address.

route --add *mgmt_subnet* **--netconfig name** [**--gateway** *default_gateway*]

Example

Configuring the default gateway IP address, 192.168.50.101, for the management subnet, 192.168.50.0/24:

```
$ route --add 192.168.50.0/24 --netconfig draft --gateway 192.168.50.101
```

Network Interfaces

Network interfaces are logical representations of physical ports on Ethernet Network Interface Cards (NICs) or Fibre Channel Host Bus Adapters (HBAs). For iSCSI traffic, each Ethernet interface must be assigned a configured subnet, and the same subnet can be assigned to multiple interfaces. You can also enable or disable VLAN tagging for each subnet on each Ethernet interface.

Fibre Channel arrays have both Ethernet and Fibre Channel interfaces. The Ethernet interfaces on a Fibre Channel array are used only for management, intra-group communication, and replication traffic; Fibre Channel interfaces are used for data traffic only. Fibre Channel interfaces do not require an assigned subnet. Instead, WWPNs are automatically assigned to them.

Add an iSCSI Interface to a Subnet

You can add (assign) an iSCSI interface to a subnet using the Draft network configuration profile to validate and activate the interface.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Add an iSCSI interface to a subnet.

   nic --assign *nic* **--netconfig draft** [**--array** {*name* | *serial*}] [**--subnet** *subnet-label*] [**--tagged** {**yes** | **no**}] [**--data_ip** *addr*]

3. Validate the Draft network configuration profile.

   netconfig --validate draft

   The array OS validates the configuration. If an error exists, the array OS returns an error message. Resolve all errors before proceeding to the next step.

4. Activate the Draft network configuration profile.

   netconfig --activate draft --force_ip_update

Example

Adding the eth1 interface to the sub5 subnet using the Draft network configuration profile:

```
$ netconfig --create_draft_from active
$ nic --assign eth1 --netconfig draft --array prod22 --subnet sub5 --tagged
yes --data_ip 192.168.50.50
$ netconfig --validate
INFO: Configuration is valid.
$ netconfig --activate
```

Remove an iSCSI Interface from a Subnet

You can remove (unassign) an iSCSI interface from a subnet using the Draft network configuration profile to validate the change. Removing an interface from a subnet disables it from carrying traffic over the subnet from which is was removed.

Procedure

1. Create a Draft network configuration profile from the Active profile.

   netconfig --create_draft_from active

2. Remove an iSCSI interface from a subnet.

   nic --unassign *nic* **--netconfig draft** [**--array** {*name* | *serial*}] [**--subnet** *subnet-label*]

   > **Note:** If an interface is added to multiple subnets, use the **--subnet** option to identify the specific subnet from which the interface is to be removed. If you do not use the **--subnet** option, the interface is removed from all subnets.

3. Validate the Draft network configuration profile.

   netconfig --validate draft

   The array OS validates the configuration. If an error exists, the array OS returns an error message. Resolve all errors before proceeding to the next step.

4. Activate the Draft network configuration profile.

   netconfig --activate draft

Example

Removing the eth1 interface from the sub5 subnet using the Draft network configuration profile:

```
$ netconfig --create_draft_from active
$ nic --unassign eth1 --netconfig draft --subnet sub5
$ netconfig --validate
INFO: Configuration is valid.
$ netconfig --activate
```

Set a Fibre Channel Interface Administrative State

You can set the administrative state of a Fibre Channel interface to either online or offline. If the Fibre Channel interface has connected initiators, you must use the **--force** option to set the administrative state to offline.

> **Note:** Using this command to disable all Fibre Channel ports on an active controller may result in pure ANO paths from the standby controller. This can lead to performance issues across all OS and single ANO path usage for IO purpose for Windows OS.

Procedure

Set the administrative state to online or offline.

fc --edit *interface_name* [**--array** {*name* | *serial*}] **--ctrlr** {**A** | **B**} **--admin_state** {**online** | **offline**} [**--force**]

Example

Setting the administrative state of the fc3 interface on controller A to online.

```
$ fc --edit fc3 --ctrlr A --admin_state online
```

Setting the administrative state of the fc1 interface on controller B to offline using the **--force** option.

```
$ fc --edit fc1 --ctrlr B --admin_state offline --force
```

iSCSI Host Connection Methods

The iSCSI initiators on the host system connect with targets on the array through the data ports on each controller. Each port is identified by its IP address. Normally, the array OS selects the IP address for each connection automatically. Hosts connect to the Virtual Target IP addresses, then the connection is automatically redirected to an appropriate iSCSI Data IP address.

The default iSCSI host connection method for the array OS releases earlier than 2.0 is *Manual*, but the default iSCSI host connection method for the array OS 2.0 and later releases is *Automatic*. However, after upgrading from a pre-2.0 release, hosts will continue to connect manually to iSCSI data IP addresses. To remedy this, install the Host Integration Toolkit on supported hosts and then change the iSCSI host connection method from Manual to Automatic. The Connection Service (CS) changes the iSCSI connections to connect to discovery IP addresses instead of data IP addresses, and NCS maintains the optimal number of connections. On the remaining hosts, change the iSCSI connections to connect to discovery IP addresses instead of data IP addresses. After all iSCSI connections are changed to connect to Virtual Target IP addresses, enable iSCSI connection rebalancing in order for the array OS to automatically rebalance iSCSI connections when distribution of connections becomes unbalanced.

If a Layer 2 inter-switch link that cannot handle the volume of iSCSI traffic is present, set up IP address zones before enabling the Automatic iSCSI host connection method. Generally speaking, a Layer 2 inter-switch link is used when traffic to different iSCSI data IP addresses goes through different switches on the same subnet.

Automatic iSCSI Host Connections

The Automatic iSCSI connection method uses data subnet discovery IP addresses for host connection. The array then automatically redirects the connection to an appropriate iSCSI data IP address. The Automatic method is the better choice for most applications.

Use a data subnet discovery IP address to connect from the host to the array.

> **Note:** The iSCSI host-to-array connection process is faster and simpler when you install and use HPE Connection Manager on your Windows or VMware host.

See the *Windows Integration Guide* and the *VMware Integration Guide*.

Manual iSCSI Host Connections

For pre-array OS 2.0 releases, hosts connect manually to iSCSI data or discovery IP addresses. The same is true for array OS 2.0 and later releases when the iSCSI connection method set to *Manual*. The Manual method is provided for legacy applications to upgrade to array OS 2.0 or later, make configuration changes, and switch to the automatic method.

Use a data subnet discovery IP address to connect from the host to the array.

> **Note:** The iSCSI host-to-array connection process is faster and simpler when you install and use HPE Connection Manager on your Windows or VMware host.

For more information, refer to the *Windows Integration Guide*, and the *VMware Integration Guide*.

Set the iSCSI Host Connection Method to Manual

The manual iSCSI host connection method is intended primarily so that you can configure legacy applications to use the automatic connection method after upgrading to release 2.0 or later. (Neither option applies to Fibre Channel arrays.)

Procedure

1. Disable automatic rebalancing of the iSCSI host connections.

   netconfig --edit active --iscsi_connection_rebalancing no

2. Disable automatic iSCSI host connections.

   netconfig --edit active --iscsi_automatic_connection_method no

Example

Example:

```
$ netconfig --edit active --iscsi_connection_rebalancing no
$ netconfig --edit active --iscsi_automatic_connection_method no
```

VLAN Support and VLAN Tagging

The array supports the configuration of VLANs, and provides a way to tag/untag frames on iSCSI or Fibre Channel arrays for specified VLANs.

> **Note:** VLANs can only be configured for use with an array after the array setup has been completed.

About VLANs

VLANs provide logical segmentation of networks by creating separate broadcast domains, which can span multiple physical network segments. They can be grouped by departments, such as *engineering* and *accounting*, or by projects, such as *release1* and *release2*. VLANs provide a number of advantages:

- Ease of administration — VLANs enable logical grouping of end-stations that are physically dispersed on a network. This aids in speed, efficiency, and accuracy of provisioning the right LUN to the right client.
- Access control — VLANs enforce security policies by separating different environments for security and compliance.
- Reduction of network traffic — By confining broadcast domains, VLANs reduce the need to have routers deployed on a network to contain broadcast traffic. In addition, end-stations on a VLAN are prevented from listening to or receiving broadcasts not intended for them. If a router is not connected between the VLANs, the end-stations of a VLAN cannot communicate with the end-stations of other VLANs.

VLANs have IDs from 1 to 4094. The array allows a single VLAN ID to be assigned to a single subnet. A subnet without a VLAN ID belongs to the default VLAN. In a group of arrays, certain restrictions apply:

- A subnet must be assigned to at least one interface on each array in the group.
- A group can have a maximum of 60 subnets, including the management subnet.
- Each array in a group can have a maximum of 120 subnet-to-NIC assignments.

Target Subnet List

Because VLANs allow you to have multiple subnets per interface, you may want to limit the number of subnets that have access to a given volume. You do this using the Target Subnet List. This is useful when there are so many subnets that timeouts may occur, for example, when subnets need to be scanned upon volume restart. It can also be used for security, to prevent certain subnets from accessing the volume.

You access the Target Subnet List when creating initiator groups.

About VLAN Tagging

Multiple VLANs can be connected through a single cable using the VLAN tagging feature. A VLAN tag is a unique identifier (between 1 and 4094) included in the 802.1Q frame header that corresponds to the ID of the VLAN to which the frame belongs. When a switch receives a tagged frame, it forwards that frame to the appropriate VLAN. A tagged frame belongs to the VLAN specified by the tag. An untagged frame belongs to the default VLAN. The array supports both tagged and untagged traffic on the same interface.

Tagged and Untagged VLAN Subnets

The VLAN tag attribute of an interface determines whether the subnet to which that interface belongs is tagged or untagged. Interfaces that have VLAN tagging enabled are exposed to the network with that subnet's VLAN ID. While a subnet can accept both tagged and untagged traffic, all tagged interface assignments must use the same VLAN ID. For each subnet, the interfaces at both ends of the network link (switch end and device end) must be assigned either tagged or untagged.

Up to 60 tagged subnets are supported on each interface, but only one untagged subnet is supported per interface. Untagged subnets are useful when running traffic over a pre-existing LAN that is not tagged, or on a switch that does not support VLAN tagging. Untagged subnets can also be used to avoid downtime when moving an array from one tagged VLAN to another.

For more information, see Move an Array from One Tagged VLAN to Another on page 45.

VLAN tagging is supported on both iSCSI and Fibre Channel subnets. For Fibre Channel, VLAN tagging is supported on management subnets only. For iSCSI, VLAN tagging is supported on both management and data subnets.

Configure VLAN Tagging

To enable VLAN tagging, first be sure a subnet has been created with a valid VLAN ID. (Subnets without a VLAN ID can only accept untagged traffic.) Then, assign network interfaces on an array to this subnet, and specify the interfaces as tagged. Use the draft configuration profile to ensure the assignment is applied consistently to all NICs using the VLAN tag. For more information on creating a draft configuration profile, refer to the *Command Line Reference*.

Note Do not use the same network for both back-end and front-end traffic. Heavy usage on one side will cause latency or congestion issues on the other side, and there will be a cascading impact to the overall environment.

Before you begin

You must already have VLANs created on your network.

Procedure

1. If no subnet exists, create a subnet with a valid VLAN ID (from 1-4094). If the subnet you want to configure already exists, skip to Step 2.

   subnet --add *subnet_name* **--subnet_addr** *netmask* **--discovery_addr** *ip_address***--type** {**mgmt** | **data** | **mgmt,data**} **--subtype** (**icsi** | **group**} **netzone_type** {**evenodd** | **bisect** | **single**} **netconfig** *name* **--vlanid** *vlan_id*

2. Create a draft configuration profile, from either the active or backup configuration profiles. You must have Power User permission or above to do this.

   netconfig --create_draft_from [**active** | **backup**]

3. Assign NICs to the draft configuration profile, with VLAN tagging set to **yes**.

   nic --assign *nic_name* **--netconfig draft --tagged yes**

   Repeat this step for each NIC for which VLAN tagging is to be added.

4. Validate the draft configuration profile.

   netconfig --validate draft

5. Activate the draft configuration profile.

   netconfig --activate draft

   Upon activation of the profile, VLAN tagging is removed from the interfaces.

6. Verify that the VLAN tagging information is correct.

   subnet --list

Example

```
$ netconfig --create_draft_from active
$ nic --assign 10GbT --netconfig draft --tagged yes
$ nic --assign 10GbE1 --netconfig draft --tagged yes
$ nic --assign 10GbE2 --netconfig draft --tagged yes
$ netconfig --validate draft
$ netconfig --activate draft
```

Remove VLAN Tagging

To remove VLAN tagging, specify the appropriate interfaces as untagged. Use the draft configuration profile to ensure the removal is applied consistently to all NICs using the VLAN tag. For more information on creating a draft configuration profile, refer to the *Command Line Reference*.

Procedure

1. Create a draft configuration profile, from either the active or backup configuration profiles. You must have Power User permission or above to do this.

   netconfig --create_draft_from [**active** | **backup**]
2. Assign NICs to the draft configuration profile, with VLAN tagging set to **no**.

   nic --assign *nic_name* **--netconfig draft --tagged no**

   Repeat this step for each NIC for which VLAN tagging is to be removed.
3. Validate the draft configuration profile.

   netconfig --validate draft
4. Activate the draft configuration profile.

   netconfig --activate draft

   Upon activation of the profile, VLAN tagging is removed from the interfaces.

Example

Removing VLAN tagging from three NICs (10GbT, 10GbE1 and 10GbE2) in the draft network configuration profile.

```
$ netconfig --create_draft_from active
$ nic --assign 10GbT --netconfig draft --tagged no
$ nic --assign 10GbE1 --netconfig draft --tagged no
$ nic --assign 10GbE2 --netconfig draft --tagged no
$ netconfig --validate draft
$ netconfig --activate draft
```

Move an Array from One Tagged VLAN to Another

The workflow below describes how to successfully move an array from one tagged VLAN to another without losing connectivity to the array. Using this method, there is always at least one path to the host, and downtime is avoided.

Procedure

1. Be sure there are at least two links between the host and the array.
2. Untag one of the links.
3. Move that link from the old host to the new host. The array is still visible on the old host using the tagged link, and is now also visible on the new host using the untagged link.
4. Untag and move the second link from the old host to the new host. Repeat this until all links have been untagged and moved.
5. Tag the links with a new VLAN ID, if desired.

Array Groups

Array groups are collections of up to four arrays that are managed as a single entity. Grouped arrays allow you to aggregate for increased performance and capacity. For iSCSI arrays, grouping makes multi-array storage pools possible.

Group Leader Array

In a multi-array group, the array serving as the group leader maintains configuration data and hosts the management IP address. In a group of iSCSI arrays, all communication received at the discovery IP address is sent to the leader of the group. Fibre Channel arrays do not use the discovery IP address.

Note: If an array is the group leader, it cannot be removed from the group. However, with the help of Nimble Storage Support, you can migrate leadership to another array in the group that stores replicated configuration data. After the leadership migration task is completed, you can remove the array that was previously the group leader.

Adding Arrays to a Group

You can add a configured (initialized) or unconfigured (uninitialized) array to an existing group.

Add an Unconfigured Array to Group

Note: You must use the group leader to add unconfigured arrays to a group.

Before you begin

- Use the same Layer 2 network as the group leader array
- Run the same version of NimbleOS as the group leader array
- Use a compatible access protocol, one of iSCSI, Fibre Channel, or multi-protocol (both iSCSI and Fibre Channel access). All arrays in a group must use the same protocol.

Procedure

1. Discover non-member arrays.

   array --discover

2. Add a discovered non-member array to the group.

   array --add *array_serial_number*

Example

```
Nimble OS $ array --discover
--------------+-------------+------------------+-----------------------------
-----
Serial Number Model         Version            Link-Local IP Addresses
--------------+-------------+------------------+-----------------------------
-----
nimble-array-1 VM-LEGACY     0.0.0.0-265702-opt 192.168.10.1, 192.168.20.2,
192.168.30.3
nimble-array-2 VM-LEGACY     0.0.0.0-262796-opt 192.168.75.50, 192.168.85.60,
```

```
 192.168.95.70
Nimble OS $ array --add nimble-array-1
Enter array name:
Enter subnet label for NIC eth1: data1
Enter subnet label for NIC eth2: data2
Enter subnet label for NIC eth3: data3
Enter subnet label for NIC eth4: mgmt-data
Enter data IP address for NIC data1(172.16.32.0/255.255.224.0): 172.16.32.1
Enter data IP address for NIC data2(172.16.32.0/255.255.224.0): 172.16.32.2
Enter data IP address for NIC data3(172.16.32.0/255.255.224.0): 172.16.32.3
eth4 is configured for management only subnet, skipping.
Enter support IP address for controller A: 172.16.224.25
Enter support IP address for controller B: 172.16.224.26
Create a new pool with this array? Enter yes or no: no
[ 1 ] default
[ 2 ] Create a new pool
Select pool to assign array [ 1 - 2 ]: 1
INFO: Array add operation may take several minutes.
```

Add a Configured Array to a Group (Group Merge)

A configured array is a group of one. Therefore, adding a configured array to a group is the same procedure as merging two groups of arrays. You can merge two groups of either Fibre Channel or iSCSI arrays. The array that is being merged into the target is the *source* group, and the target array is the *destination* group.

> **Note:** There is no way to "unmerge" groups. To undo a merge, you must remove the arrays from the merged group. You can then create a new group for those arrays, if desired.

Before you merge two groups, verify the following requirements.

Platform Compatibility Notes:

- You cannot merge an iSCSI array and a Fibre Channel (FC) array together in the same group.
- You can add multiple FC arrays to the same group.
  - NimbleOS 3.x.x.x or later is required.
  - Fibre Channel WWPNs will change for the array added to the group. You must adjust zoning appropriately to reflect this
- All flash (AFA), Hybrid Flash (HFA), and Secondary Flash (SFA) arrays can be merged into the same multi array group.
  - Pools of different array types (AFA, HFA, and SFA) cannot be merged into a single multi-array pool.
  - Only pools from the same array type can be merged into a multi-array pool.
- You can add an AFA to an existing HFA or SFA group, migrate to the AFA using volume move, and remove the old array non-disruptively.
  - NimbleOS 4.x.x.x or later is required.
  - For more information, see KB-000277
- You can add an HFA or SFA to an existing AFA group, migrate from the HFA or SFA to AFA, and remove the HFA or SFA array.
  - This process is disruptive to host I/O since this iSCSI arrays will undergo a discovery IP change.
- If the same CHAP username is configured on both the source and destination groups, only the CHAP user on the destination array will be used after the merge.
- If secure shell (SSH) keys exist for users of the same name on source and destination groups, the SSH keys for the destination group will be used. SSH keys for that user in the source group will be discarded.

Verify the following:

- Both groups use the same data and management subnets along with the same subnet labels.

  - All array groups to be merged must have the same subnets defined using the same subnet names
  - All interfaces assigned to a given subnet must be able to reach each other, regardless of which array they are located on.
  - All interfaces within a given subnet must be within the same broadcast domain (the L2 or layer 2 network segment).
  - All interfaces in a particular subnet must be in the same VLAN, and any inter-switch links that exist between those interfaces must be configured to allow the same VLAN
  - The array will have one or more subnets tagged as Allow Group. For any subnet where Allow Group is set to Yes, the switch ports connected to those arrays need to also be in the same native VLAN.
  - No more than 4 data subnets can exist.
- Both groups use the same data protocol (iSCSI or Fibre Channel) and have the data ports connected.
- All arrays to be merged that have Active Directory integration enabled must be configured to use the same Active Directory domain for authentication.
- Both arrays are in an active or standby state.
- Both arrays have the same NimbleOS version.
  - If not, update one or both arrays. To learn more about updating the NimbleOS, see the *HPE Hardware Guide*.
- Throttle levels are the same for both groups.
  - Group-level throttles from the joining group are automatically discarded and the hosting group throttles are automatically applied. If the throttle levels are set differently on the two arrays that are associated with the hosting and joining groups, adjust the throttle level on the hosting group based on your environment and requirements.
- Arrays must have the same maximum transmission unit (MTU) setting on configured subnets.
- Switches and inter-switch links must have their MTU's set to a value equal to or greater than the MTU used on group members.
  - Note that some switches require that the MTU be set higher than the value on the array, due to differences in how they calculate the value of the MTU.
  - There are no consequences to setting the switch MTU to the largest reasonable value, often as high as 9214 or 9216 bytes.

Notes on Names:

- All array names in your environment must be unique.
- iSCSI subnet label names on both arrays are the same.
- Group names are unique and are short enough to avoid truncation when merged.
- The group name is different from the name of the array being added.
- Volume names cannot be re-used on the source and destination groups.
- All initiator group names must be unique.
  - Initiator IQNs on the source and destination, if matching, also use the same case.
- All protection schedule names must be unique.
- The names for all new pools must contain fewer than 64 characters, if you intend to merge two groups.
  - If the name of the source group (sometimes referred to as group B) is short enough to form a new pool name that is less than 64 characters, then truncation does not occur. However, if the source group has a long name, such as a 63-character name, then the new pool name results in more than 64 characters. In this case, the name is automatically truncated to the limit of 64 characters.
- Other than the default pool, pool names on the source and destination groups must be unique.
  - The default pool on the source group will be renamed `default-source_group_name`.
  - A group named `default-source_group_name` cannot exist on the destination group.

Pre-Merge Tasks

Note:

- Array OS 4.X is required.
- Merging an iSCSI array and a FC array together in the same group is not supported.
- You cannot perform a pool merge between different array types, for example AFA, HFA and SFA.
  - To add an AFA to an existing HFA or SFA group and keep both (with their own pools), you must migrate to the AFA and remove the old array non-disruptively.

    For more information, see *KB-000277*.
  - To add an HFA or SFA to an existing AFA group, you must migrate from the HFA or SFA to AFA, and remove the HFA or SFA array.
  - This process is disruptive to host I/O and not recommended for iSCSI arrays which will undergo a discovery IP change.
- Deduplication Domains cannot be merged, see *Domains* in the *Deduplication* section for details.

Verify the following factors:

- Both arrays are in an active or standby state and are running the same array OS version.

  If not, update one or both arrays. To learn more about updating the array OS, refer to the Installation Guide or Hardware Guide for your array model.
- Switches and inter-switch links must have their MTU's set to a value.
- One or more subnets must be set to include group traffic in the **Traffic Assignment** menu. For any subnet where Allow Group is set to Yes, the switch ports connected to those array interfaces must be in the same Native VLAN.Arrays must have the same MTU setting on configured subnets.
- Verify that the following factors are the same for both groups:
  - The data and management subnet
  - The data protocol

    The data ports must be connected.

    Note:

    A data subnet is required so the management subnet must be modified for use for management and data. The management and data ports must be in the same VLAN (L2 network, broadcast domain) as the other interfaces in that subnet.

    Data IPs must be added for group traffic between the arrays.
  - The L2 network or broadcast domain

    Both groups must also be in the same VLAN, on the same switch or switch-stack, or on switches connected via an inter-switch link or trunk.
  - The replication throttle levels

    Group-level throttles from the source group are automatically discarded and the destination group throttles are automatically applied.

    If the throttle levels are set differently on the two arrays that are associated with the destination and source groups, adjust the throttle level on the destination group based on your environment and requirements.

Complete the following tasks if applicable:

- Configure source and destination groups to have matching subnets in the Active network configuration.

  - Delete any draft network configs from source and destination groups.
  - On source and destination groups, “Edit” and “Save” the network configuration in order to over-write the “backup” network configuration.
- Configure the joining group (the group being merged) as a single-array configuration.
  - You can merge only two groups at a time. Make sure that the joining group consists of a single array.
- Perform the following replication tasks if you intend to merge replication partners into the same group:
  - Stop all replication and group merge processes in your environment.
  - Remove the hosting and joining groups from any replication configurations.
- Unregister the HPE vCenter plugin, if you previously registered it.
  - To learn more about unregistering the plugin, see the *VMware Integration Guide*.
- Set offline all volumes associated with the joining group.

  Note: It is recommended that you perform this task using the GUI.

For Fibre Channel arrays:

- After the volumes are offline, then offline all Fibre Channel interfaces on the source group. These tasks are easier using the GUI.
- A data subnet is still required
  - The management subnet can be modified to be used for Mgmt+Data
  - Alternately, an unused Ethernet interface can be configured as a “data” interface, and a data IP added for group traffic between the arrays.
- An alias is a human-friendly name associated with a WWPN. An alias may be defined on the switch (as a fabric-assigned alias) or using the array OS (as a user-assigned alias). For any two disjoint groups with user-assigned aliases, two potentially problematic situations can occur:
  - Different initiators with the same WWPN must use the same alias for that WWPN
  - Different initiators cannot be using the same alias for different WWPNs
  - In either case, you are presented with this information on the validation screen, and must manually perform the rename operation.

Merge Two Groups

Before you begin

- Check array cabling. For more information about cabling, refer to the *Nimble Storage Quick Start Guide* for your arrays.
- Ensure that the network configuration is correct.

Procedure

1. Resolve conflicts before merging groups.

   group --merge_validate *group_name* [**--username** *username*] [**--password** *password*]

2. Merge the two groups.

   group --merge *group_name* [**--username** *username*] [**--password** *password*]

   Note: The merge command also runs the validation.

Example

Validating and merging two groups.

```
Nimble OS $ group --merge_validate MyGroup --username MyName --password MyPass▶
word
Nimble OS $ group --merge MyGroup --username MyName --password MyPassword
```

Post-Merge Tasks

After you merge two groups, perform the following tasks.

Procedure

1. Online all volumes that were previously in the source group.
2. Update any SNMP clients to contact the group leader's (destination group's) management IP rather than the source group's, and the group leader's SNMP community name (rather than the source group's).
3. For iSCSI groups, update initiators with the new discovery IP for these volumes and then reconnect.
4. Re-arrange pool and volume assignments as needed. For example, merge the default source group into the default pool on the destination group.
5. If there was replication going to the source group, change it to now go to the destination group.
   a) On any group (other than the destination group) that was previously replicating to the source group, create a partner to the destination group.
   b) Edit all schedules previously using the source group so that the destination group is used instead.
   c) Delete the source group as a partner.
6. If there was replication going from the source group, change it to now come from the destination group.
   a) On any group (other than the destination group) that was previously being replicated from the source group, create a partner to the destination group.
   b) Promote all volume collections that were replicated from the source group (making the new group the owner).
   c) Demote to the destination group.

      This makes the destination group the new owner for these volume collections.

   d) On the destination group, all partners that were paused on the source group before the merge must be resumed.
7. Rezone.

   If you had zoning set up, then re-zone due to changes to WWPNs for what was previously the source group.

Group Information

When you select **Hardware** in the GUI, a summary of the arrays in a group and a summary of the group itself appears. For single-array groups, the Hardware Details page opens.

The following information for each array is displayed, as well as the total aggregate for the group:

Item	Description
Group and Arrays	The name of the group and all member arrays
IOPS	The IOPS measured in 5-minute increments
MiB/s	The performance as measured in 5-minute increments
Usage	Array usage and group capacity
Used (capacity)	Amount of capacity used for each array and for the group
Storage Pool, if applicable	Storage pools to which each member is assigned and totals

Default Group Settings

Because you manage group members as a single entity, you perform many administrative actions for the group instead of an individual arrays. You can modify the following default group settings:

- Password for the Admin user
- Period of inactivity before an array logs users out
- Default space reservation
- Date, time, and time zone
- DNS settings

Note: You must have an Administrator role to make these settings.

Modify the Default Inactivity Timeout

By default, users are logged out after a specified amount of time with no activity. You can change the default inactivity timeout value up to a maximum of 720 minutes (12 hours). The initial default value is 30 minutes.

CAUTION: Before you change the default inactivity timeout interval, consider all security issues.

Before you begin

You must have Administrator permission to change the default inactivity timeout interval for the group.

Procedure

1. Change the global user inactivity timeout (in minutes).

   group --edit --inactivity_timeout *minutes*

2. Verify the new user inactivity timeout

   group --info

Example

Configuring the global user inactivity timeout to 240 minutes.

```
$ group --edit --inactivity_timeout 240
$ group -info
Group name: test-group-1
  .
  .
  .
User inactivity timeout: 240 minute(s)
```

Modify the Default Space Reservation

You can set the default space reservations for new volumes or replicas. If needed, you can specify different settings when creating volumes.

Procedure

Modify the default volume reserve percentage.

group --edit default_vol_reserve *percent*

Example

Modifying the default volume reserve percentage.

```
Nimble OS $ group --edit default_vol_reserve 75
```

Modify the Date, Time, and Time Zone

The date and time are set when the array is created. However, you can change the date, time, and time zone at any time.

Procedure

1. Modify the time zone for an NTP server

   group --edit --ntpserver *server*

   timezone --set *zone*

2. Modify the time, date, and time zone for an array.

   date [**--utc**] [**--edit** {*hh:mm[ss]* |*'YYYY-MM-DD hh:mm[:ss]'*}]

   timezone --set *zone*

   When you use the --utc option, change the Coordinated Universal Time.

Example

Modifying the date and time for an array.

```
$ date --edit 2015-09-30 12:00:00
```

Modifying the time zone for an NTP server.

```
$ group --edit --ntpserver 10.25.55.155
$ timezone --set Europe/Stockholm
```

Modify DNS Settings

You can use a DNS server to map a hostname to an IP address. The service enables users to type host names that can be translated into an IP address usable by networking software. You can modify the DNS server settings for the group as needed.

Procedure

Modify the group's DNS settings.

group --edit [**--dnsserver** *server*] [**--domainname** *domain name*]

Example

Modifying the DNS server hostname or IP address.

```
$ group --edit --dnsserver 10.25.50.255
```

Modifying the DNS to allow support to establish a connection to the group to collect diagnostic information.

```
$ group --edit --support_tunnel yes
```

Initiator Groups

An initiator is a port on a server or computer that "initiates" a connection with "target" ports on a storage array. It can be an Ethernet Network Interface Card (NIC) port that initiates a connection over an iSCSI fabric to one or more target ports, or a Fibre Channel Host Bus Adapter (HBA) port that initiates a connection over a Fibre Channel fabric to one or more target ports.

An initiator group is a collection of either iSCSI initiators or Fibre Channel initiators that are managed as a single unit.

iSCSI Initiator Groups

An iSCSI initiator group is a collection of one or more iSCSI initiators, with each initiator having a unique iSCSI Qualified Name (IQN) and IP address. Each IQN represents a single Network Interface Card (NIC) port on an iSCSI-based client in the form of a Windows server, ESXi or Linux host. Configure iSCSI initiator groups on the array; configure client-side iSCSI initiators according to the vendor's recommendations.

> **Note:** By default, iSCSI volumes deny access to initiators. To allow initiators in an initiator group to access a volume, you must configure an ACL that includes the desired initiators and attach it to the volume.

Create an iSCSI Initiator Group

Before you begin

To create an iSCSI initiator group, you must know either the iSCSI Qualified Name (IQN) or IP address, or both for each initiator being added to the group.

Procedure

1. Create an initiator group.

   initiatorgrp --create *initiator_group_name* [**--description** *text*]

2. Add an initiator to the group.

   initiatorgrp --add_initiators *initiator_group_name* [**--label** *initiator_label*] [**--initiator_name** *IQN*] [**--ipaddr** *ip-addr*] [**--initiator_alias** *alias*]

   > **Note:** If you cannot copy-and-paste the IQN, type it very carefully.

3. Add a subnet.

   initiatorgrp --add_subnets *initiator_group_name* [**--label** *subnet_label*]

4. Verify your iSCSI initiator group configuration.

   initiatorgrp --info *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --create Datamon
Nimble OS $ initiatorgrp --add_initiators Datamon --label Basic --initiator_name
 iqn.1991-05.com.microsoft:techops.storage.com --ipaddr 192.0.2.88
Nimble OS $ initiatorgrp --add_subnets Datamon --label Subnet-198.51.100.0
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: iscsi
Created: Jan 01 2016 12:00:01
```

```
Last Configuration Change: Jan 01 2016 12:00:01
Number of Subnets: 1
        Subnet Label: Subnet-198.51.100.0
Number of Initiators: 1
        Initiator Label: Basic
                Initiator Name: iqn.1991-05.com.microsoft:techops.storage.com
                Initiator IP Address: 198.0.2.88
Number of Volumes: 2
        Volume Name: northVol
        Volume Name: southVol
```

Edit an iSCSI Initiator Group

Procedure

1. Modify the initiator group name, description, or host type.

   initiatorgrp --edit *initiatorgrp_name* [**--name** *new_name*] [**--description** *text*] [**--host_type** *host_type*]

2. Add a subnet to the iSCSI initiator group.

   initiatorgrp --add_subnets *initiatorgrp_name* [**--label** *subnet_label*]

3. Remove a subnet from the iSCSI initiator group.

   initiatorgrp --remove_subnets *initiatorgrp_name* [**--label** *subnet_label*]

4. (Optional) Verify your iSCSI initiator group configuration.

   initiatorgrp --info *initiatorgrp_name*

Example

```
$ initiatorgrp --edit Datamon
$ initiatorgrp --name Dataman
$ initiatorgrp --description datamanager
$ initiatorgrp --info Dataman

Name: Dataman
Description:
Access Protocol: iscsi
Created: Jan 01 2016 12:00:01
Last Configuration Change: Jan 01 2016 12:00:01
Number of Subnets: 1
        Subnet Label: Subnet-198.51.100.0
Number of Initiators: 1
        Initiator Label: Basic
                Initiator Name: iqn.1991-05.com.microsoft:techops.storage.com
                Initiator IP Address: 198.0.2.88
Number of Volumes: 2
        Volume Name: northVol
        Volume Name: southVol
```

Delete an iSCSI Initiator Group

Before you begin

Before you can delete an iSCSI initiator group, you must remove all associated volumes from the initiator group's access control list(s) (ACLs). For more information, see Remove an Initiator Group ACL from a Volume on page 61.

Procedure

1. Verifly that there are no volumes associated with the iSCSI initiator group.

   initiatorgrp --info *initiator_group_name*

If any volumes are listed for the initiator group, remove them. For more information, see Remove an Initiator Group ACL from a Volume on page 61.

2. Delete the initiator group.

   initiatorgrp --remove *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: iscsi
Created: Jan 01 2016 12:00:01
Last Configuration Change: Jan 01 2016 12:00:01
Number of Subnets: 1
        Subnet Label: Subnet-198.51.100.0
Number of Initiators: 1
        Initiator Label: Basic
                Initiator Name: iqn.1991-05.com.microsoft:techops.storage.com
                Initiator IP Address: 198.0.2.88
Number of Volumes: 0
Nimble OS $ initiatorgrp --delete Datamon
```

Add an Initiator to an iSCSI Initiator Group

Procedure

1. Add an initiator to the iSCSI initiator group.

   initiatorgrp --add_initiators *initiator_group_name* [**--label** *initiator_label*] {**--initiator_name** *IQN* | **--ipaddr** *client_system_IP-address*}

   Note: If you cannot copy-and-paste the IQN, type it very carefully.

2. Verify that the initiator has been added to the iSCSI initiator group.

   initiatorgrp --info *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --add_initiators Datamon --label Intermediate --ini▶
tiator_name iqn.1991-06.com.microsoft:techops.storage.com --ipaddr 192.0.2.90
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: iscsi
Created: Jan 01 2016 12:05:01
Last Configuration Change: Jan 01 2016 12:05:01
Number of Subnets: 1
        Subnet Label: Subnet-198.51.100.0
Number of Initiators: 2
        Initiator Label: Basic
                Initiator Name: iqn.1991-05.com.microsoft:techops.storage.com
                Initiator IP Address: 198.0.2.88
        Initiator Label: Intermediate
                Initiator Name: iqn.1991-06.com.microsoft:techops.storage.com
                Initiator IP Address: 198.0.2.90
Number of Volumes: 4
        Volume Name: northVol
        Volume Name: southVol
        Volume Name: eastVol
        Volume Name: westVol
```

Remove an Initiator from an iSCSI Initiator Group

Procedure

1. Remove an initiator from the iSCSI initiator group.

   initiatorgrp --remove_initiator *initiator_group_name* [**--label** *initiator_label*]

2. Verify that the initiator has been removed from the iSCSI initiator group.

   initiatorgrp --info *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --remove_initiator GST-Test --label Basic
```

```
Nimble OS $ initiatorgrp --info GST-Test
Name: GST-Test
Description:
Host Type: auto
Access Protocol: iscsi
Application identifier:
Created: Sep  4 2018 14:01:50
Last configuration change: Oct  1 2020 13:46:08
Number of Subnets: All
Number of Initiators: 0
Number of Volumes: 2
        Volume Name: Vol2ForTesting
        Volume Name: bks-volume
```

Fibre Channel Initiator Groups

A Fibre Channel initiator group is a collection of one or more initiators, with each initiator having a unique World Wide Port Name (WWPN). Each WWPN represents a single Host Bus Adapter (HBA) port on a Fibre Channel-based client in the form of a Windows server, ESXi or Linux host. Configure Fibre Channel initiator groups on the Nimble array; configure client-side Fibre Channel initiators according to the vendor's recommendations.

> **Note:** By default, Fibre Channel volumes deny access to initiators. To allow initiators in an initiator group to access a volume, you must configure an ACL that includes the desired initiators and attach it to the volume.

Create a Fibre Channel Initiator Group

Before you begin

To create an Fibre Channel initiator group, you must know the Fibre Channel World Wide Port Name (WWPN) for each initiator being added to the group.

Procedure

1. Create a Fibre Channel initiator group.

   initiatorgrp --create *initiator_group_name* [**--description** *text*] [**--host_type** *host_type*]

   > **Note:** The --host_type option is only needed when creating an initiator group for HP-UX systems. Specify hpux as the host_type sub-option when creating an initiator group in HP-UX.

2. Add an initiator.

   initiatorgrp --add_initiators *initiator_group_name* [**--label** *label*] [**--initiator_alias** *initiator_alias*] **--wwpn** *WWPN*

   Example:

Note: A WWPN is 16 hexadecimal characters (case insensitive) in XX:XX:XX:XX:XX:XX:XX:XX or XXXXXXXXXXXXXXXX format. If you cannot copy-and-paste the WWPN, type it very carefully.

3. Verify your Fibre Channel initiator group.

   initiatorgrp --info *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --create Datamon
Nimble OS $ initiatorgrp --add_initiators Datamon --label client1 --wwpn
10:20:00:a0:fa:63:c2:61
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: fc
Created: Jan 01 2016 12:00:01
Last Configuration Change: Jan 01 2016 12:00:01
Number of Initiators: 1
        Initiator: client1 (10:20:00:a0:fa:63:c2:61)
Number of Volumes: 1
        Volume Name: northVol
                LUN: 0
```

Edit a Fibre Channel Initiator Group

Procedure

1. Change the initiator group name or description.

   initiatorgrp --edit *initiator_group_name* [**--name** *new_name*] [**--description** *text*] [**--host_type** *host_type*]

   Note: The --host_type option is only needed when editing an initiator group for HP-UX systems. Specify hpux as the host_type sub-option when editing an initiator group in HP-UX.

2. Add an initiator.

   initiatorgrp --add_initiators *initiator_group_name* [**--label** *label*] [**--initiator_alias** *initiator_alias*] **--wwpn** *WWPN*

   Example:

   Note: A WWPN is 16 hexadecimal characters (case insensitive) in XX:XX:XX:XX:XX:XX:XX:XX or XXXXXXXXXXXXXXXX format. If you cannot copy-and-paste the WWPN, type it very carefully.

3. Remove an initiator from the initiator group.

   initiatorgrp --remove_initiators *initiator_group_name* [**--label** *initiator_label*] **--wwpn** *WWPN*

   Note: If you cannot copy-and-paste the WWPN, type it very carefully.

4. Verify your Fibre Channel initiator group configuration.

   initiatorgrp --info *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --create Datamon
Nimble OS $ initiatorgrp --add_initiators Datamon --label client1 --wwpn
10:20:00:a0:fa:63:c2:61
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
```

```
Access Protocol: fc
Created: Jan 01 2016 12:00:01
Last Configuration Change: Jan 01 2016 12:00:01
Number of Initiators: 1
        Initiator: client1 (10:20:00:a0:fa:63:c2:61)
Number of Volumes: 1
        Volume Name: northVol
                LUN: 0
```

Delete a Fibre Channel Initiator Group

Before you begin

Before you can delete an Fibre Channel initiator group, you must remove all associated volumes from the initiator group's access control list(s) (ACLs). For more information, see Remove an Initiator Group ACL from a Volume on page 61.

Procedure

1. Verifly that there are no volumes associated with the Fibre Channel initiator group.

   initiatorgrp --info *initiator_group_name*

   If any volumes are listed for the initiator group, remove them. For more information, see Remove an Initiator Group ACL from a Volume on page 61.

2. Delete the initiator group.

   initiatorgrp --remove *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: fc
Created: Jan 01 2016 12:00:01
Last Configuration Change: Jan 01 2016 12:00:01
Number of Initiators: 1
        Initiator: client1 (10:20:00:a0:fa:63:c2:61)
Number of Volumes: 0
Nimble OS $ initiatorgrp --delete Datamon
```

Add an Initiator to a Fibre Channel Initiator Group

Procedure

1. Add an initiator.

   initiatorgrp --add_initiators *initiator_group_name* [**--label** *label*] [**--initiator_alias** *initiator_alias*] **--wwpn** *WWPN*

   Example:

   Note: A WWPN is 16 hexadecimal characters (case insensitive) in XX:XX:XX:XX:XX:XX:XX:XX or XXXXXXXXXXXXXXXX format. If you cannot copy-and-paste the WWPN, type it very carefully.

2. Verify that the initiator has been added to the iSCSI initiator group.

   initiatorgrp --info *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --create Datamon
Nimble OS $ initiatorgrp --add_initiators Datamon --label client1 --wwpn
```

```
10:20:00:a0:fa:63:c2:61
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: fc
Created: Jan 01 2016 12:00:01
Last Configuration Change: Jan 01 2016 12:00:01
Number of Initiators: 1
        Initiator: client1 (10:20:00:a0:fa:63:c2:61)
Number of Volumes: 1
        Volume Name: northVol
                LUN: 0
```

Remove an Initiator from a Fibre Channel Initiator Group

Procedure

1. Remove an initiator from the initiator group.

   initiatorgrp --remove_initiators *initiator_group_name* [**--label** *initiator_label*] **--wwpn** *WWPN*

   Note: If you cannot copy-and-paste the WWPN, type it very carefully.

2. Verify that the initiator has been removed from the Fibre Channel initiator group.

   initiatorgrp --info *initiator_group_name*

Example

```
Nimble OS $ initiatorgrp --create Datamon
Nimble OS $ initiatorgrp --add_initiators Datamon --label client1 --wwpn
10:20:00:a0:fa:63:c2:61
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: fc
Created: Jan 01 2016 12:00:01
Last Configuration Change: Jan 01 2016 12:00:01
Number of Initiators: 1
        Initiator: client1 (10:20:00:a0:fa:63:c2:61)
Number of Volumes: 1
        Volume Name: northVol
                LUN: 0
```

Initiator Group Access Control Lists

All initiators in an initiator group are granted access to a volume when the access-control list (ACL) for an initiator group is added to the volume. An ACL can be added to multiple volumes, granting the initiators in the group access to those volumes.

When you create or edit a volume, you can add one or more initiator group ACLs to it.

Note: If you do not add any ACLs to a volume, no initiators will be able to connect to the volume.

Add an Initiator Group ACL to a Volume

Procedure

1. Add an initiator group access control list (ACL) to a volume.

   vol --addacl *volume_name* [**--apply_acl_to** *volume* | *snapshot* | *both*} **--initiator_group** *group_name*

2. Verify that the initiator group ACL has been added to the volume.

   initiatorgrp --info *group_name*

Example

```
Nimble OS $ vol --addacl test2 --apply_acl_to volume --initiator_group Datamon
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: fc
Application identifier:
Created: Feb 10 2016 12:00:01
Last configuration change: Jan 01 2016 12:00:01
Number of Initiators: 3
        Initiator: winInit2 (12:34:56:78:90:12:34:52)
        Initiator: esxiInit22 (12:34:56:78:90:12:34:54)
        Initiator: linuxInit78 (12:34:56:78:90:12:34:56)
Number of Volumes: 2
        Volume Name: test1
                LUN: 0
        Volume Name: test2
                LUN: 1
```

Remove an Initiator Group ACL from a Volume

Procedure

1. Remove a volume from an iSCSI Initiator Group Access Control List (ACL)

   vol --removeacl [*volume_name*] **--apply_acl_to** [*volume* | *snapshot* | *both*] **--initiator_group** [*group_name*]

2. Verify that the initiator group ACL has been removed from the volume.

   initiatorgrp --info *group_name*

Example

```
Nimble OS $ vol --removeacl test1 --apply_acl_to volume --initiator_group
Datamon
Nimble OS $ initiatorgrp --info Datamon
Name: Datamon
Description:
Access Protocol: fc
Application identifier:
Created: Feb 10 2016 12:00:01
Last configuration change: Jan 01 2016 12:00:01
Number of Initiators: 3
        Initiator: winInit2 (12:34:56:78:90:12:34:52)
        Initiator: esxiInit22 (12:34:56:78:90:12:34:54)
        Initiator: linuxInit78 (12:34:56:78:90:12:34:56)
Number of Volumes: 1
        Volume Name: test2
                LUN: 1
```

Volumes

Volumes are the basic storage units from which the total capacity of an array is apportioned. The number of volumes per array depends on how the storage is allocated.

Hosts connect to volumes using iSCSI or Fibre Channel. A volume appears to a host as a single disk drive, which can be used as a file system, a raw disk, or a virtual disk.

When you delete a volume, the snapshots that are associated with that volume are also deleted. If the volume has online snapshots, they must be taken offline before you can delete them.

Clones, Replicas, and Snapshots

The array OS lets you manage the objects in a storage system: volumes and their associated clones, replicas, and snapshots.

Clones are writable, highly space-efficient copies of volumes which you can create from snapshots. When you create a clone from a snapshot, you create a new volume with a new name and iSCSI or Fibre Channel target with the same settings. Clones share identical blocks and are often used to test applications before putting them into production.

Replicas are copies of volumes stored on a different array, called a replication partner. Replicas are most often used for disaster recovery. For more information about replicas and replication, see Replication on page 104.

Snapshots are point-in-time copies of volumes. Snapshots are often used as backups, and to preserve the state of volumes at specific points. By creating a clone from a snapshot, snapshots can also be used as starting points to which applications can write and read data. For more information about snapshots, see Snapshots on page 97.

Logical versus Physical Space

When working with volumes, it is important to understand the difference between logical space and physical space.

Physical storage resources are aggregated into storage pools from which the logical storage is created. It allows you to have a logical space for data storage on physical storage disks by mapping space to the physical location. Physical space is the actual space on the hardware that is used. For example, when you set a volume or snapshot reserve, that physical space is reserved and taken out of the general pool of space. It is physical resource consumption.

Logical space is space that the system manages, such as the volume size. In this case, the volume size is not necessarily the actual amount of space on a physical disk, but the amount of space defined for a volume, which may span multiple physical disks.

Space Management

The array OS has built-in capacity saving mechanisms such as inline compression and thin provisioning. The following considerations help you plan your space configuration for volumes and snapshots.

The simplest form of space management is to not use reserves at all. This means that there is no dedicated (prereserved) space per volume taken from the general storage pool, so all volumes can consume what they need as it is needed. This method requires that you monitor space usage to ensure that there is always space available.

However, for critical volumes, such as those hosting business-critical data, it may be more important for you to reserve space to ensure that the volume will always have enough. Reserved space is immediately taken from the storage pool.

When you create a volume, you define a certain amount of space for that volume. The volume space is the size that is reported to your application.

Volume Reserve

The volume reserve is the guaranteed minimum amount of physical space reserved for the volume. Reserved space is set aside for the volume and is immediately withdrawn from the general storage pool. Set the volume reserve from 100% (the entire physical space is reserved) to 0% (no physical space is reserved). As new data is written to the volume, the free space within the volume reserve decreases.You can increase the volume reserve if needed.

One consideration when setting volume sizes and reserves is the level of compression you get for a particular application or data set. For example, most volumes should see 50-75% compression, so a volume reserve set to 10 GB will be able to store far more than the actual 10 GB space if it were uncompressed. In other words, 10 GB space of application data will only use between 2.5 GB and 5 GB when compressed.

Note: Setting the reserve to 100% effectively disables thin provisioning.

Thin Provisioning

Thin provisioning is a storage virtualization technology that uses physical storage only when data is written instead of traditional provisioning, which reserves all the capacity up-front when an application is configured. This method addresses over-provisioning and its associated costs. Frequently, volumes reserve excessive space against expected growth. Often this growth does not materialize, or materializes much later than expected. With thin provisioning, you create volumes and assign them to servers and applications, but the physical resources are only assigned when the data is written. Physical storage that is not being used remains available to other volumes. No unnecessary storage is reserved for use by any single application.

For example, like most SANs, your array must support several applications. Projections show that eventually the total storage needed by all applications will reach 3 TB. However, for the first few quarters of the year, these applications should only use about 300 GB. Instead of creating the volumes using the total 3 TB that you expect to need, with thin provisioning you can create three 1 TB volumes, but set the reserve to only 150 GB for each volume. When you factor in compression savings, the applications should not use the full 3 TB until the next purchasing window, minimizing the cost of buying more capacity until it is needed.

Volume Usage Limits

Volume usage limits determine how much of a volume can be consumed before an alert is sent to the administrator. When the usage limit is reached, the performance policy associated with the volume determines the next action (for example, whether to make the volume read-only or take the volume offline). An alert is also sent.

You can configure warning levels (a volume usage warning level and a snapshot usage warning level) below the usage limits so you get an alert before the limit is actually reached.

In most cases, you should set the quota to 100% (no quota). Some applications do not tolerate changes to volume sizes. Quotas were developed to address this issue. Using quotas lets you set a limit but leave room in case more space is needed.

For example, if you have an application that you do not want to fill all the space on the volume before more space is available for expansion, set a quota for the volume.You now have a safety factor, and when the quota is met, you can reset the quota, giving more space to the application.You can then plan for further expansion if necessary.

If the volume is approaching the quota limit, an event is logged. If enforcement is enabled, the administrator can access the system log to determine what follow-up actions to take, such as preventing the user from accessing more disk space or allocating additional disk space to the user.

Note: Volume usage limit must be greater than or equal to both the volume reserve and the volume warning level.

A Note on Defragmentation

Do not defragment volumes on an array. The value of defragmentation is mainly on a local physical disk to keep files contiguous so the disk heads do not require unnecessary physical seeks across the platters, and slow down file I/O.

There is no such value about the effectiveness of this in a networked iSCSI environment, especially where files are stored on storage devices that have their own layers of virtualization.

Defragmenting files in a storage array environment results in changed blocks, even though the files did not change, and can have unnecessary impacts, such as snapshots being larger than they should.

Cloning Space Considerations

Clones are space-efficient copies of a volume that can be used independently of the source volume. When created, they have the same settings as the volume from which they were created. Clones share blocks that are identical with the source volume, and only begin to use space when changes are made.

HPE recommends that you lower the reserve settings for clones. When determining the reserve settings, factors to consider include how long-lived the clone will be and how much the clone will vary from its source. For example, the reserve settings may not need to be very high if the clone is being run to test a new application against, will not be changed much, and will be deleted after the testing is complete.

Note: You cannot delete the source volume of a clone unless you first delete the clone.

Protecting Data Using Snapshots

Snapshots ensure that data stored in volumes is always recoverable.

Note: If a volume is deleted, all associated snapshots are also permanently deleted.

Merging primary and backup storage makes snapshots an efficient method to protect data. Because no data needs to be copied outside the array, snapshots can be created and used to restore data almost instantaneously.

You can restore a volume from either the local recovery point or the remote recovery point. The local recovery point is the last local snapshot taken for the volume. The remote recovery point is the last snapshot taken through replication.

Because snapshots are part of the converged storage and backup, and because they are so efficient, consider the implications when creating snapshots. For some applications, the amount of storage used for snapshots may equal or exceed the storage needed for the source volume. Reserving storage for snapshots does not reserve the space from the space allocated to the volume, but from the general storage pool.

Even if you plan to manually take snapshots of volumes or use a third-party program to create backups, create a volume collection without schedules on volumes that are being manually snapshotted.

Create a Volume

Volumes are also referred to as logical units (LUNs). They are the building blocks of any storage system. A host connects to the volume via iSCSI or Fibre Channel, and the volume appears to the host as a single disk drive, which can be used as a file system, a raw disk, or a virtual disk.

Note:

The--chapuser and --multi_initiator options for the vol command apply only to iSCSI volumes.

The --cached_pinned and --dedupe_enabled options cannot both be enabled on the same volume created on a hybrid flash array that supports deduplication.

Procedure

1. Create a volume.

   vol --create*name* **--size** *megabytes* [**--description** *text*] [**--perfpolicy** *name*] [**--iops_limit** *iops*] [**--mbps_limit** *mbps*] [**--cached_pinned** {*yes* |*no*}] [**--dedupe_enabled** {*yes* |*no*}] [**--reserve** *percent*] [**--quota** *percent*] [**--warn_level** *percent*] [**--snap_reserve** *percent*] [**--snap_quota** *percent*] [**--snap_warn_level** *percent*] [**--start_offline**] [**--apply_acl_to** {*volume*

|*snapshot* |*both*}] [**--chapuser** *username*] [**--initiatorgrp** *group name*] [**--lun** *lun*] [**--multi_initiator** {*yes* |*no*}] [**--pool** *pool name*] [**--folder** *folder name*] [**--agent_type** {*smis* |*none*}] [**--encryption_cypher** {*aes-256-xts* |*none*}]

Important:

If the volume description contains spaces, enclose all characters and spaces of the *text* variable in quotation marks. If you do not, the volume description will be truncated at the first space.

See *vol* in the *Command Reference* for more detail about this command.

Note: If deduplication is enabled when creating the volume, the volume and snapshot reserves are set to 0. Once deduplication is enabled for a volume, the volume reserve, snapshot reserve, and application category cannot be changed, even if deduplication is later disabled.

2. (Optional) Add an ACL to grant access to the volume to an initiator group.

   vol --addacl *test1* **--initiatorgrp** *group_name*

   Note:

   See Create an iSCSI Initiator Group on page 54 for the procedure to create an iSCSI initiator group.

   See Create a Fibre Channel Initiator Group on page 57 for the procedure to create a Fibre Channel initiator group.

Example

Creating a volume named test1 and specifying its size.

```
Nimble OS $ vol --create test1 --size 10000
```

Creating a volume named test1, specifying its size, and setting the encryption cypher to AES 256 XTS.

```
Nimble OS $ vol --create test1 --size 10000 --encryption aes-256-xts
```

Creating a volume named test1, specifying its size, providing a description, and setting percent values for quota and warning level.

```
Nimble OS $ vol --create test1 --size 10000 --description "for all departments"
 --quota 80 --warn_level 75
```

Creating a volume named test1, specifying its size, applying the performance policy Policy1, and setting the cache pinning option to yes.

```
Nimble OS $ vol --create test1 --size 10000 --perfpolicy Policy1 --cached_pinned
 yes
```

Creating a volume named test1, specifying its size, and setting percentages for snapshot reserve, snapshot quota, and snapshot warning level.

```
Nimble OS $ vol --create test1 --size 10000 --snap_reserve 10 --snap_quota 10
 --snap_warn_level 80
```

Adding an ACL to grant access to the volume test1 to the initiator group Dataman.

```
Nimble OS $ vol --addacl test1 --initiatorgrp Dataman
```

What to do next

- Configure the iSCSI or Fibre Channel connections on your volumes to connect to the group leader array.
- Configure the iSCSI or Fibre Channel connections for your server or host to access those volumes.

- Configure your client initiator according to the vendor's recommendations.

Edit a Volume

You can modify most volume configuration options, such as access restrictions, quotas, performance policy, and volume collection assignment after you have created the volume.

Note: If deduplication was enabled when the volume was created, the volume reserve, snapshot reserve, and application category cannot be changed.

Before you begin

Before you rename a volume, you must take it offline.

Procedure

1. Take the volume offline.

   vol --offline *vol_name*

2. Make the required changes.

   vol --edit *vol_name* [**--name** *new_name*] [**--size** *megabytes*] [**--description** *text*] [**--perfpolicy** *name*] [**--iops_limit** *iops*] [**--mbps_limit** *mbps*] [**--cached_pinned** {*yes* |*no*}] [**--dedupe_enabled** {*yes*| *no*}] [**--readonly** {*yes* | *no*}] [**--force**] [**--reserve** *percent*] [**--quota** *percent*] [**--warn_level** *percent*] [**--snap_reserve** *percent*] [**--snap_quota** *percent*] [**--snap_warn_level** *percent*] [**--multi_initiator** {*yes* | *no*}] [**--agent_type** {*none* | *smis*}]

   Important:

   If the volume description contains spaces, enclose all characters and spaces of the *text* variable in quotation marks. If you do not, the volume description will be truncated at the first space.

   See *vol* in the *Command Reference* for more information about this command.

   Note: If you are editing a volume with reserve, and you enable deduplication, the volume and snapshot reserve will both be set to 0.

3. Set the updated volume to online.

   --online *vol_name*

4. (Optional) Add an ACL to grant access to the volume to an initiator group.

   vol --addacl *test1* **--initiatorgrp** *group_name*

   Note:

   See Create an iSCSI Initiator Group on page 54 for the procedure to create an iSCSI initiator group.

   See Create a Fibre Channel Initiator Group on page 57 for the procedure to create a Fibre Channel initiator group.

Example

Taking the volume test1 offline and changing its name to test2.

```
Nimble OS $ vol --offline test1
Nimble OS $ vol --edit test1 --name test2
```

Setting the volume test2 online.

```
Nimble OS $ vol test2 --online
```

Editing the volume test2, changing its size, description, and the percent values for quota and warning level.

```
Nimble OS $ vol --edit test2 --size 20000 --description "for finance department"
 --quota 85 --warn_level 80
```

Editing the volume test2, changing the performance policy, and setting the cache pinning option

```
Nimble OS $ vol --edit test2 --perfpolicy Policy3 --cached_pinned no
```

Adding an ACL to grant access to the volume test2 to the initiator group Dataman.

```
Nimble OS $ vol --addacl test2 --initiatorgrp Dataman
```

Change a Volume's State

Taking a volume offline makes that volume unavailable to initiators. When you set a volume to offline, all current connections are closed.

Procedure

Set a volume offline or online, depending on its current state. Use one of the following commands:

- **vol --offline** *vol_name*
- **vol --offline** *vol_name*

Example

Setting a volume offline:

```
$ vol --offline MyVolume
```

Setting a volume online:

```
$ vol --online MyVolume
```

Clone a Volume

Cloning a volume (via a snapshot) creates a new volume with a new name, but keeps all other settings of the original, including the data at the time the snapshot was taken.

Clones of snapshots are useful for restoring individual files instead of a complete volume. By copying the files from the cloned volume to the active volume, you can restore the files without affecting other users.

When you clone a volume, settings such as reported size, reserve size, quotas, snapshot reserve size, and security are cloned. Volume collections are not automatically assigned to the clone. Clones are set online by default, and are writable. The clone consumes space from the same space as the original volume.

> **Note:** Only the mandatory options required to create a clone using the vol command are given. For information about all options that can be used to create a clone, see the *vol* entry in the *Command Reference*.

Procedure

Clone a volume.

vol --clone *vol_name* **--snapname** *snap_name* **--clonename** *clone_name*

Example

Creating the clone Clonetest1 from volume Test1 and snapshot Snap30.

```
Nimble OS $ vol --clone Test1 --snapname Snap30 --clonename Clonetest1
```

What to do next

Edit the clone to assign it to the desired volume collection to ensure that snapshots and replicas will be made according to the desired schedule.

Restore a Volume from a Snapshot

You can restore a volume from one of its snapshots. A snapshot is automatically taken of the existing state before the volume is restored, even if third-party software is used.

Note: When restoring a volume, it is recommended that you unmount the volume from the host before putting the restored snapshot online. Restoring a volume without stopping all host access can cause data corruption and system errors.

Procedure

1. Set the volume to be restored offline.

   vol --offline *vol-name*

2. Restore the volume from the specific snapshot you need to restore the data.

   vol --restore *vol-name* **--snapname** *name*

3. Set the restored volume online.

   vol --online *vol_name*

Example

Restoring the volume nimVol3 from the snapshot nimSnap12:

```
$ vol --offline nimVol3
$ vol --restore nimVol3 --snapname nimSnap12
$ vol --online nimVol3
```

Delete a Volume

Deleting a volume also deletes any snapshots of the volume. Because clones share the original data with the source volume, you cannot delete a volume that has a clone.

CAUTION: All data stored on the volume will be destroyed.

Procedure

1. Take the volume offline.

   vol --offline *vol_name*

2. Disassociate the volume from any volume collections.

   vol --dissoc *vol_name*

   Note: This step is mandatory for all volumes, including standalone volumes.

3. Delete the volume.

   vol --delete *vol_name*

Volume Pinning

Volume pinning allows you to keep active blocks of a volume in the cache, as well as writing them to disk. This provides a 100% cache hit rate for specific volumes (for example, volumes dedicated to critical applications), and delivers the response times of an all-flash storage system.

A volume is "pinned" when the entire active volume is placed in cache; associated snapshot (inactive) blocks are not pinned. All incoming data after that point is pinned. The number of volumes that can be pinned is limited by the size of the volumes and amount of available cache. However, only one volume in a volume family (a volume and the associated snapshots and clones) can be pinned.

> **Note:** Pinning a volume may affect the cache hit rate of other volumes. It is a best practice to avoid unnecessary cache pinning.

Pinnable Flash Capacity

Only a portion of the array's total flash capacity can be used for pinning. The amount of pinnable capacity is determined by the amount of usable cache in the system and the amount of usable disk capacity in the system. The array OS performs these calculations.

The formula for determining pinnable capacity depends on whether you are using a deduplication:

- A hybrid array without deduplication:

  pinnable capacity = 66% * (usable cache capacity - (4% * usable disk capacity))

  For example, an array without deduplication enabled that has a usable disk capacity of 100 TB and a total flash capacity of 12 TB, the total pinnable capacity of the array is 66% of ((12 TB - (4% * 100 TB)) = 5.28 TB.

- A hybrid array with deduplication enabled:

  pinnable capacity = 66% * (usable cache capacity - (4% * usable disk capacity) – (4% * maximum-enabled deduplication capacity))

For striped pools, the pinnable capacity is

(minimum value of (pinnable capacity/usable disk capacity) on any individual array) * (total usable disk capacity across all arrays in the stripe)

If you are unable to pin a volume, you may need to adjust the flash capacity. For more information, see Unable to Pin a Volume on page 71.

Volume Pinning Caveats

When performing certain operations related to volume pinning, keep these caveats in mind.

Table 8: Volume Pinning Caveats

Condition	Solution
Volume Pinning Enablement	To be able to pin the volume, the cache must be enabled on the performance policy for the volume.
Quota and Volume Size Changes	When changing a pinned volume's quota, the amount of cache reserved for the volume will also change. You will not be able to pin volumes until sufficient cache is available.

Condition	Solution
Moving a Volume	When moving a volume, the pinnable capacity on the destination array or pool must be sufficient to pin the moving volume. Before moving the volume, ensure there is enough cache on the destination array or pool. The cache reservation on the source will be freed on the completion of the move. If you attempt to pin a volume during a volume move, pinning will not be guaranteed until a scan has finished on the destination. This could take some time (up to several hours) depending on how much content is in the volume, and how much other activity there is.
Replicating a Volume	If you replicate a pinned volume, the volume will not be pinned downstream.
Adding Shelf Capacity	When adding a shelf to an array, if the array's new pinnable capacity is less than the size of the pinned volumes, the shelf activation will fail.
Failed, Removed, or Upgraded SSD	When an SSD fails or is removed, if the amount of pinnable capacity is less than the current size of the pinned volumes, then all volumes may become unpinned because there is not enough free cache to pin the blocks. A message is displayed recommending that you consider adding capacity or unpinning some volumes to restore performance to cache pinned volumes. If the SSD is replaced with one of the same capacity, volume pinning continues normally. If the SSD is replaced with one of a different capacity, the amount of pinnable cache is recalculated. One consolidated alert is sent for all volumes when the free usable cache drops below the acceptable level, and another alert is sent once the free usable cache returns to an acceptable level for the rescan to be completed.
Pinning an Existing Volume	If you are using the Edit function to specify that a volume is to be pinned, this initiates a scan of the volume. Pinning will not be guaranteed until the scan of the volume is finished. This could take some time (up to several hours) depending on how much content is in the volume, and how much other activity there is. Alerts are sent both when pinning begins and ends, on a per-volume basis.
Performing a Bin Migration or Pool Merge	If you have performed a bin migration (for space balancing, for example), or if you attempt to pin a volume during a bin migration or pool merge, pinning will not be guaranteed until a rescan has been completed on the bin's new destination. This could take some time (up to several hours) depending on how much content is in the volume, and how much other activity there is. For a bin migration, one consolidated alert is sent for all volumes once the rescan is finished. For a pool merge, one consolidated alert is sent for all the volumes before and after the merge.
Performing a Volume Snapshot Restore	If you want to do a volume snapshot restore (unpin an old volume and pin a new volume) to re-establish the heat map (cache hit information), pinning will not be guaranteed until a scan of the new tip has finished. This could take some time (up to several hours) depending on how much content is in the volume, and how much other activity there is. **Note:** A snapshot restore can only be performed after unpinning the volume.
Performing a Software Upgrade	In case of a software upgrade, pinning will not be guaranteed until a rescan is performed to determine whether any blocks that were in memory were evicted as a result of the software upgrade. This could take some time (up to several hours) depending on how much content is in the volume, and how much other activity there is. One consolidated alert is sent for all the volumes regarding possible loss of pinning, and another alert is sent once the rescan has completed.

Condition	Solution
Unmanaged Shutdown	In the case of an unmanaged shutdown (where a few segments of pinned data are not flushed and can show up as cache misses), pinning will not be guaranteed until a rescan of the system is finished. One consolidated alert is sent for all the volumes before and after the rescan.
Volume Promotion/Demotion	Demoting a pinned volume will automatically unpin the volume. The volume can be pinned again after promoting it.

Pinning a Volume Created with a Previous Release

When trying to pin an existing volume that was created before version 2.1, the volume will not be pinned. This is because the current array caching technology is not compatible with volumes created before Version 2.1.

To successfully pin a volume that was created before array OS version 2.1, you must upgrade the array to 2.1 or higher, and wait for the cache to age out. As the cache ages out under the new array OS, it is rebuilt with the latest caching technology. Once the cache ages out completely, you can safely pin the volume. A two-month cache aging period is generally sufficient. However, if you want to obtain the exact cache aging period for your volume, contact support.

Unable to Pin a Volume

If you try to pin a volume and receive a "cache capacity exceeds available capacity" message, you have insufficient usable cache to pin the volume. There are two options:

- Use another pool - you can use the array OS to display a list of other pools with sufficient cache. You have the option to switch your volume to that pool.
- Unpin other pinned volumes - you can view the pinning capacity of other volumes by hovering over their names listed on the Caching dialog. A tooltip is displayed with the pinning capacity information for that volume. If your volume is already pinned but you want to free more usable cache, you can use the Caching facet on the volume's information page to view a list of pinned volumes.

Pin a Volume

Note: Before you pin a volume, the performance policy associated with that volume must have caching enabled.

Procedure

Pin the volume.

vol --edit *volume_name* **--cache_pinned yes**

If the volume can be pinned, you will see additional capacity information for that volume. If the volume cannot be pinned, you may see a "cache capacity exceeds available capacity" message.

Example

Pinning a volume:

```
vol --edit volume1 --cache_pinned yes
```

Unpin a Volume

You can unpin an existing volume using the commands below.

Procedure

Unpin a volume.

vol --edit *volume_name* **--cache_pinned no**

When a volume is unpinned, the blocks are given a cache performance policy of Normal.

Example

Unpinning a volume:

```
vol --edit volume1 --cache_pinned no
```

Performance Policies

Performance policies define how data is stored on the array to achieve optimal performance for a specific application. There are several predefined performance policies to choose from based on the application assigned to the volume.

Note: The default performance policy for a Secondary Flash array is Backup Repository.

Create a Performance Policy

Important: When you replicate a volume using a performance policy, use an identical policy for the volume on the replication partner.

Procedure

Create a performance policy.

perfpolicy --create *name* [**--description** *text*] [**--blocksize** *bytes*] [**--compress** {**yes** | **no**}] [**--cache** {**yes** | **no**}] [**--cache_policy** {**normal** | **aggressive**}] [**--space_policy** {**offline** | **non_writable**}] [**--app_category** *category*]

Note:

By default, a performance policy has no description, a block size of 4096 bytes, compression and cache are enabled, cache policy is normal, and space policy is offline. If you want to set other values, change those options when you create the performance policy.

The --app_category (application category) option is case sensitive. For more information about the perfpolicy command, see the *Command Reference*.

Example

Creating a performance policy named Hunter with a block size of 8192 bytes, an aggressive cache policy, and an application category of Exchange:

```
$ perfpolicy --create Hunter --blocksize 8192 --cache_policy aggressive --
app_category Exchange
```

Create a Performance Policy with Deduplication Enabled

Note: Because deduplication can be enabled only on All Flash, Secondary Flash, and some models of Adaptive Flash arrays, this task applies only to those arrays that are not striped in a single pool.

Procedure

Use the **yes** argument of the **--dedupe_enabled** option to create a performance policy with deduplication enabled.

perfpolicy --create *name* [**--description** *text*] [**--blocksize** *bytes*] [**--compress** {**yes**|**no**}] [**--cache** {**yes**|**no**}] [**--cache_policy** {**normal**|**aggressive**}] [**--space_policy** {**offline**|**non_writable**}] [**--app_category** *category*] [**--dedupe_enabled** {**yes**|**no**}]

Edit a Performance Policy

Note: You cannot edit or delete any of the predefined performance policies.

Procedure

Edit at least one performance policy parameter.

perfpolicy --edit *name* [**--name** *new_name*] [**--description** *text*] [**--compress** {**yes** | **no**}] [**--cache** {**yes** | **no**}] [**--cache_policy** {**normal** | **aggressive**}] [**--space_policy** {**offline** | **non_writable**}] [**--app_category** *category*]

Example

Editing the description and cache parameters for the performance policy named 64kBlocksize:

```
$ perfpolicy --edit 64kBlocksize --description "64k block size performance
policy"
--cache no
```

Delete a Performance Policy

Note: Only performance policies that are not currently associated with a volume can be deleted.

Procedure

1. (Optional) View a list of the performance policies, and identify the one to be deleted.

   perfpolicy --list

2. Delete the performance policy.

   perfpolicy --delete *name* [**--force**]

   Note: Use the **--force** option only if you must delete a performance policy associated with a volume or folder.

Example

Deleting the performance policy named 64kBlocksize:

```
$ perfpolicy --list
---------------------------------+----------+--------+-----+------------
Performance Policy               Block Size Compress Cache Cache Policy
Name                             (bytes)
---------------------------------+----------+--------+-----+------------
64kBlocksize                          65536 Yes      Yes   aggressive
default                                4096 Yes      Yes   normal
Exchange 2003 data store               4096 Yes      Yes   normal
Exchange 2007 data store               8192 Yes      Yes   normal
Exchange 2010 data store              32768 Yes      Yes   normal
Exchange log                          16384 Yes      No    normal
Hyper-V 2012 VDI Storage               4096 Yes      Yes   normal
Hyper-V 2012 VM Storage                4096 Yes      Yes   normal
Hyper-V CSV                            4096 Yes      Yes   normal
Oracle OLTP                            8192 Yes      Yes   normal
Other                                  4096 Yes      Yes   normal
SharePoint                             8192 Yes      Yes   normal
SQL Server                             8192 Yes      Yes   normal
SQL Server 2012                        8192 Yes      Yes   normal
SQL Server Logs                        4096 Yes      No    normal
VMware ESX                             4096 Yes      Yes   normal
VMware ESX 5                           4096 Yes      Yes   normal
```

```
VMware VDI                                4096 Yes      Yes   normal
vSphere Datastore for Exchange            4096 Yes      Yes   normal
vSphere Datastore for SQL Server          4096 Yes      Yes   normal
VVol Operating System                     4096 Yes      Yes   normal
VVol VDI                                  4096 Yes      Yes   normal
Windows File Server                       4096 Yes      Yes   normal
$ perfpolicy --delete 64kBlocksize
```

Virtual Volumes (vVols)

Virtual volumes is VMware's functionality for management of VMs and their data (such as VMDKs and physical disks). The ability to manage virtual volumes using VMware virtual disks mapped to volumes is a new feature in vSphere 6.0.

Both vVols and regular volumes can coexist on the same array or set of arrays (group or pool). vVols are visible in both the CLI and GUI as regular volumes, for monitoring their capacity and performance. However, you use the vCenter UI to manage vVols. For more information, see the *VMware Integration Guide*.

When you edit, delete, offline, or online a vVol using the **vol** commands, a warning message is displayed.

You can perform limited tasks in the array OS, as described in this section.

vCenter Server

VMware vCenter Server is a data center management server application developed by VMware Inc. to monitor virtualized environments. vCenter Server provides centralized management and operation, resource provisioning and performance evaluation of virtual machines residing on a distributed virtual data center. A vCenter Server is used to manage the VVols feature in NimbleOS.

A vCenter *extension* must be registered with the vCenter server before it can be displayed in the arrays that connect to the vCenter server.

You can register or unregister extensions, and add, edit, or remove a vCenter Server using NimbleOS.

Register a vCenter Plugin with vCenter Server

To manage volumes or Virtual Volumes (vVols) through vCenter Server, you must register the HPE vCenter Plugin with vCenter Server.

Procedure

1. Add a vCenter Server.

   vcenter --add [**--name**] [**--hostname** {*host_name*|*ip_addr*}] [--port_number *port_number*] [**--username** *user_name*] [**--password** *password*] [**--description** *description*] [**--subnet_label** *subnet_label*]

2. Register a vCenter Plugin with vCenter Server.

   vcenter --register *vCenter_name* [**--extension** {**web**| **thick**| **vasa**}]

   Note: To manage vVols, specify the **--extension vasa** command option to register the vCenter extension for a VASA provider.

Unregister a vCenter Plugin

If you no longer want to manage volumes or Virtual Volumes (vVols) through vCenter Server, you can unregister the HPE vCenter Plugin currently registered with vCenter Server.

Procedure

Unregister the vCenter Plugin.

vcenter --unregister *vCenter_name* [**--extension** {**web** | **thick** | **vasa**}]

Add a vCenter Server

To manage volumes through vCenter, you must add a vCenter Server to the array.

Procedure

Add a vCenter Server.

vcenter --add [**--name**] [**--hostname** {*host_name*|*ip_addr*}] [--port_number *port_number*] [**--username** *user_name*] [**--password** *password*] [**--description** *description*] [**--subnet_label** *subnet_label*]

Edit a vCenter Server

You can edit the parameters for a vCenter Server that has already been added to an array.

Procedure

Edit a vCenter Server.

vcenter --edit *vCenter_name* [**--name** *vCenter_name*] [**--username** *username*] [**--password** *password*] [**--description** *description*]

Remove a vCenter Server

You can remove a vCenter Server that has already been added to an array.

Procedure

Remove a vCenter Server.

vcenter --remove *vCenter_name*

Virtual Machines

Before Virtual Volumes, a deleted VM could be easily restored from a snapshot. However, with vVols, the delete occurs directly on the volume. Therefore, HPE has provided a "deferred deletion" functionality that allows you to restore or permanently delete VMs that have been used with vVols. The configuration vVol is preserved only if there is at least one unmanaged snapshot that can be used to restore from. The data vVol is preserved even if no snapshots have been taken.

Viewing Accidentally Deleted VMs

You can view a list of VMs or any volumes related to VMs that are deleted, before you permanently delete them. Add a section here that says "Accidentally deleted VMs". Command "vm --list --deleted" will show the list of VMs or any volumes related to VMs that are deleted.

Procedure

View accidentally deleted VMs.

vm --list --deleted

Permanently Delete a VM

This command will permanently delete vVols which have been deferred-deleted (deleted from the vCenter). This command will not delete volumes which are still in use and have not been deleted from the vCenter. Use this command only if you are sure the VM will not need to be restored.

Note: You must have Power user privileges or higher to run this command.

Procedure

Permanently delete a VM.

vm --destroy [**--name** *vmname*] [**--id** *vmid*]

Restore a VM

This command restores all volumes associated with a virtual machine that was deleted accidentally. If the VM name is not unique, you must specify the VM ID.

The vCenter admin is then responsible for browsing the vVol datastore and adding the VM back to the inventory.

Note: You must have Power user privileges or higher to run this command.

Procedure

Restore a virtual machine.

vm --restore [**--name** *vmname*] [**--id** *vmid*]

Folders

Folders are containers for holding volumes. They are used most often for organization, management, and further delegation. Folders provide simple volume grouping for ease of management.

You can monitor the performance of folders by going to the **Monitor** > **Performance** tab.

External management agents such as VASA Storage Containers (Nimble Virtual Volumes) and SMI-S storage pools, map to folders and can leverage them directly.

A folder can have a usage limit, a provisioned limit, or no limit.

- Usage limit – Limits the amount of space used by volumes and clones in the folder. The usage includes the compressed size of both the volumes and the snapshots. For virtual volumes, the usage limit is the size that will be reported for the datastore in vCenter. Note that the usage limit will cause new volume creation to fail when the usage limit is reached.
- Provisioned limit – Limits the amount of space that can be provisioned in the folder.

Relationship of Folders, Pools, and Volumes

It is important to know the characteristics of folders and volumes and their relationship to each other and to pools. Some of the characteristics are outlined in the following table.

Folders	Volumes
Folders are provisioned within pools, and can contain volumes.	A volume can belong to a pool without being part of a folder.
Folder names must be unique within the pool that contains them.	Volume names, even those within folders, must be unique across a group.
Pools containing folders can be merged after name conflicts are resolved. However, the folders themselves cannot be merged.	Volumes and their clones can be spread across multiple folders. Volumes can be moved across folders in the same pool or in multiple pools.

Create a Folder

Use the **folder --create** command to create a folder, specify a space limit, and folder agent_type. Folder names can be up to 64 characters, and up to 64 folders can be created per group. Folders cannot be created at the vSphere cluster level.

> **Note:** If the agent type is VVol, a VASA Provider must be registered before folder creation, and an appserver must be specified using the **folder --create** command. You cannot specify an appserver for other agent types. If the agent type is SMI-S, a performance policy must be specified in the **folder --create** command. You cannot specify a performance policy for other agent types.

Procedure

Create a folder.

folder --create *folder_name* [**--pool** *pool_name*] [**--description** *text*] [**--usage_limit** *mebibytes*] [**--agent_type** {**none**|**smis**|**vvol**|**openstack**}] [**--appserver** *vcenter_name*] [**--perfpolicy** *perfpolicy*] [**--iops_limit** *iops*] [**--mbps_limit** *mbps*]

Example

Creating the folder finance in the default pool, with an application server of vcenter1, agent type of vvol, and a usage limit of 10000 mebibytes.

```
folder --create finance --pool default --appserver vcenter1 --agent_type vvol
 --usage_limit 10000
```

Edit a Folder

Procedure

Edit a folder.

folder --edit *folder_name* [**--pool** *pool_name*] [**--name** *name*] [**--description** *text*] [**--usage_limit** *mebibytes*] [**--appserver** *vcenter_name*] [**--force**] [**--perfpolicy** *perfpolicy*] [**--iops_limit** *iops*] [**--mbps_limit** *mbps*]

Example

Rename a folder.

```
folder --edit folder_name --name name
```

Delete a Folder

A folder can be deleted only if it is empty (contains no volumes).

Procedure

Delete a folder.

folder --delete *folder_name* [**--pool** *pool_name*]

> **Note:** You cannot delete a folder with volumes in it. You must first move or delete the volumes.

Deduplication

Deduplication is a form of data reduction that saves storage space. The deduplication process identifies duplicate content within a domain, and stores only one copy of that content.

The HPE storage array implementation of deduplication works at the volume block level on the following arrays:

- All Flash arrays running release 3.x or later
- Secondary Flash arrays running release 4.2.0 or later
- Select models of Adaptive Flash arrays running release 5.0.1 or later

When deduplication is enabled, identical content stored on the array is deduplicated using inline deduplication. Inline deduplication involves arrays deduplicating data in real-time, as data is received.

The deduplication process uses a two-level fingerprint system, with short fingerprints for speed of detection and long cryptographically secure fingerprints to ensure reliability. The deduplication process optimizes for “flocks” of duplicate data, consecutive runs of blocks that are duplicated. This multi-layer deduplication process allows for near-perfect duplication detection, while dramatically reducing the amount of main memory required to efficiently deduplicate large capacity SSDs.

If data has already been written to the array with deduplication disabled, the data on the disk cannot be deduplicated unless you migrate the data either using array-side functionality (for example, move the volume to another deduplication-enabled pool in the group) or host tools to migrate to a new deduplication-enabled volume or pool.

Note: Volume (and snapshot) limits and reserves are based on pre-deduplication usage.

Deduplication on Hybrid Arrays

Deduplication is supported on the following hybrid arrays:

- CS500, CS700
- CS1000, CS3000, CS5000, CS7000
- HF20, HF20H, HF40, HF60

These restrictions apply to deduplication on hybrid arrays:

- Pinned volumes cannot be deduplicated.
- Deduplication cannot be enabled across striped pools.
- Replicated data is not deduplicated; the data is replicated without any deduplication savings.
- The array must include the number of SSD drives indicated in the following table:

Array Model	Required Number of SSDs
HF20H	2 SSDs
HF20H fully populated	4 SSDs
HF20H fully populated and upgraded to HF40H	4 SSDs
HF20, HF40, HF60	6 SSDs
CS500, CS700	4 SSDs
CS1000	3 SSDs
CS3000, CS5000, CS7000	6 SSDs

Note: See the *Array Configuration Matrix* available on HPE InfoSight at https://infosight.hpe.com/ for more information.

Note: If a hybrid platform contains volumes with deduplication enabled, any extra flash capacity that results from unpinning a volume is used to increase deduplication capacity.

Note:

Arrays that are updated from release 5.0.2.0 and 5.0.1.0 might have volumes with deduplication enabled. Any arrays that are updated to release 5.0.3.0 or later with deduplicated volumes will operate as a deduplication capable array, regardless of the number of installed SSDs. Such configurations are *not* recommended by HPE.

The following tables provide information about the Maximum Deduplication Capacity (MDC) on supported hybrid arrays and the additional Flash to Disk Ratio (FDR) required to support MDC.

MDC and pool deduplication capacity outputs apply to hybrid arrays in the CS series only. On HF series and later models, the entire array capacity can be deduplicated.

Note:

You must have a four percent FDR to enable deduplication on the hybrid models. For MDC, you must have an additional four percent FDR for a total of eight percent FDR.

To see the deduplication capacity (TiB), log in to the array OS CLI as an administrator and run the **pool --info** *pool_name* command. On the HF20H, HF20, HF40, and HF60 models, this command returns **N/A** as the value for dedupe capacity (MiB). This is because you can enable deduplication for the entire array.

Table 9: Effective Capacity and Additional Flash Required to Support MDC

Platform	Maximum Deduplication Capacity (MDC)	Effective Capacity with 3x Deduplication	Additional Flash Required to Support MDC
CS500	40 TiB	120 TiB	1.6 TiB
CS700	100 TiB	300 TiB	4 TiB
CS1000	10 TiB	30 TiB	0.4 TiB
CS3000	40 TiB	120 TiB	1.6 TiB
CS5000	100 TiB	300 TiB	4 TiB
CS7000	200 TiB	600 TiB	8 TiB

Note: Before you enable deduplication on hybrid arrays, review the product documentation for complete details.

Pool-Level Deduplication (Default)

The default setting varies depending on the array type. These are the default settings:

- For pools consisting of a single All Flash array, and Secondary Flash array or Adaptive Flash array (HFxx): The pool-wide deduplication capability is turned on. All newly created volumes in the pool are created with deduplication enabled. If you turn the setting off, all newly created volumes in this pool inherit the deduplication setting defined by their performance policy.
- For pools consisting of a single Adaptive Flash array (CSxxxx, CSxxx): The pool-level deduplication setting is not available. All newly created volumes in this pool inherit the deduplication setting defined by their performance policy.

Domains

A deduplication domain is the area where deduplication takes place. It is defined by three settings: intersection of containers (same pool, or folder hierarchy), performance policy, and block size.

Volumes that share blocks are grouped together in the same deduplication domain. Two characteristics help determine which volumes can be grouped together: application category and block size. An application category is an attribute indicating that the volumes store data from the same type of application. Application categories are predefined, and cannot be changed. They are selected when creating or updating a performance policy. Volumes with the same block size are able to be deduplicated, and can be part of the same deduplication domain. Volumes with different block sizes, performance policies or pools cannot be part of the same deduplication domain.

Cloned volumes inherit the parent deduplication domain; you cannot create a clone in another domain. To move a clone to another domain, you use the volume move operation.

Deduplication domains cannot be merged. Also, deduplication is not supported where the volumes are striped across a pool. Pools containing deduplicated volumes cannot be merged with other pools. You cannot add an array to a pool containing an All Flash array that has volumes with deduplicated blocks.

Enable All-Volume (Pool-Level) Deduplication

When creating or editing a pool, use the **--dedupe_all_volumes yes** option to enable deduplication by default on all newly created volumes. Existing volumes are not affected. Deduplication on existing volumes can be changed by editing the volume attributes.

Note: You cannot enable pool-level deduplication when creating or editing pools on adaptive flash arrays.

Procedure

pool --create *pool_name* **--array** *array* [**--description** *description*] [**--dedupe_all_volumes yes**]

Enable Deduplication Determined by Performance Policy

When creating or editing a pool, use the **--dedupe_all_volumes no** option to disable deduplication by default on all newly created volumes. Note that existing volumes are not affected.

When set to **no**, the deduplication setting for the volume is inherited from the pool's performance policy. When new volumes are created, they inherit the deduplication setting defined in the performance policy. Deduplication on existing volumes can be changed by editing the volume attributes.

Note: Pool-level deduplication is not applicable to pools consisting of adaptive flash arrays.

Procedure

pool --edit *pool_name* [**--dedupe_all_volumes no**]

Enable Per-Volume Deduplication

Procedure

Enable deduplication on a single volume.

vol --edit *name* **--size** *mebibytes* [**--description** *text*] [**--perfpolicy** *name*] [**--dedupe_enabled** {**yes**|**no**}]

Example

Enable volume deduplication.

```
vol --edit vol1 --size 60 --description TestVol --perfpolicy policy1 --
dedupe_enabled yes
```

Create volume with deduplication enabled.

```
vol --create vol1 --size 60 --description TestVol --perfpolicy policy1 --
dedupe_enabled yes
```

Disable Per-Volume Deduplication

Procedure

Disable deduplication on a volume.

vol --edit *name* **--size** *mebibytes* [**--description** *text*] [**--perfpolicy** *name*] [**--dedupe_enabled** {**yes**|**no**}]

Example

Disable per-volume deduplication

```
vol --edit vol1 --size 60 --description TestVol --perfpolicy policy1 --
dedupe_enabled no
```

Create a volume with deduplication disabled.

```
vol --create vol1 --size 60 --description TestVol --perfpolicy policy1 --
dedupe_enabled no
```

Create a Performance Policy with Deduplication Enabled

Note: Because deduplication can be enabled only on All Flash, Secondary Flash, and some models of Adaptive Flash arrays, this task applies only to those arrays that are not striped in a single pool.

Procedure

Use the **yes** argument of the **--dedupe_enabled** option to create a performance policy with deduplication enabled.

perfpolicy --create *name* [**--description** *text*] [**--blocksize** *bytes*] [**--compress** {**yes**|**no**}] [**--cache** {**yes**|**no**}] [**--cache_policy** {**normal**|**aggressive**}] [**--space_policy** {**offline**|**non_writable**}] [**--app_category** *category*] [**--dedupe_enabled** {**yes**|**no**}]

Clone a Volume with Deduplication Enabled

Procedure

Use the **yes** argument of the **--dedupe_enabled** option to clone a volume with deduplication enabled.

vol --clone *vol_name* **--snapname** *snap_name* **--clonename** *clone_name* [**--description** *text*] [**--readonly** {**yes**|**no**}] [**--reserve** *percent*] [**--quota** *percent*] [**--warn_level** *percent*] [**--snap_reserve** *percent*] [**--snap_quota** *percent*] [**--snap_warn_level** *percent*] [**--start_offline**] [**--apply_acl_to** {**volume**|**snapshot**|**both**}] [**--chapuser** *user_name*] [**--initiatorgrp** *group_name*] [**--lun** *lun*] [**--multi_initiator** {**yes**|**no**}] [**--cache_pinned** {**yes**|**no**}] [**--dedupe_enabled** {**yes**|**no**}]

Storage Pools

You can divide the storage in multi-array groups into multi-array storage pools. For example, you might have a segregated storage pool for certain applications, users, or workloads. Arrays can be members of only one pool at a time. However, volumes assigned to storage pools can span multiple arrays.

The difference between groups and storage pools is that groups aggregate arrays for management, while storage pools aggregate arrays for capacity and performance. Storage pools provide automatic load balancing and migration capability, if all the other prerequisites for pools are met.

Pool Considerations

In appropriate environments, you can add or remove group members from a storage pool or you can combine storage pools.

A single-array storage pool provides fault isolation. Volumes whose pools are on one array keep data placement simple and snapshots truly local. All host operating systems identified in the *Validated Configuration Matrix* and array data protocols support single-array pools.

A multi-array pool lets you consolidate capacity and scale performance. Volumes striped across multiple array pools can have larger capacity, pool-wide resources, and more even growth with less rebalancing later. Some host operating systems support multi-array pools.

Consider the following points when planning whether and how to implement storage pools.

- Operating system on host machines

  Windows, ESX, and Linux hosts for which a supported Connection Manager is available can use multi-array storage pools. Any hosts for which a supported connection manager is unavailable require single-array pools.

- Disk speed
- Disk capacity
- Network bandwidth
- Application using the storage pool
- Load balancing requirements
- Volumes that will reside on the storage pool

If you are planning on pooling arrays across dissimilar models, consider the following caveats.

- Coupling a lower performance model array with a higher performance model could degrade the performance of the pooled system.
- Performance can be affected when you merge arrays into a single pool when the arrays have different performance caps, cache configurations, and storage capacity.
- Consider using a multi-array pool to migrate data from an older array to a newer array, as this method can be less disruptive than using replication to migrate data.

Storage pools can include some or all of the arrays in a group. When you add an array to a group, you must also add this array to one of the existing storage pools for that group or you can create a new storage pool and add the array to the new pool.

The first implementation example shows a group that includes three HPE arrays that are all in the same group in the default pool.

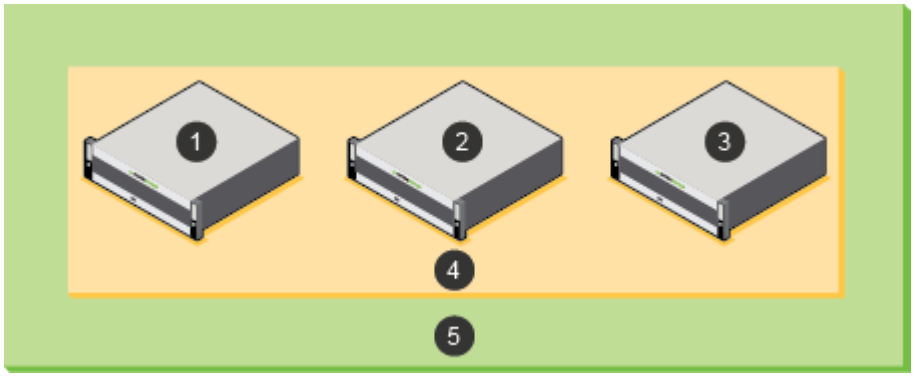

1 Array 1

2 Array 2

3 Array 3

4 Default pool

5 Multi-array group that includes arrays 1, 2, and 3

The second implementation example shows two pools that allow a specific application, such as Microsoft Exchange, to use volumes that are spread across two of three arrays in the group, but not use the storage on a third array. In this scenario, the storage on the array labeled 3 is in the default pool. To configure this example, remove the arrays labeled 1 and 5 from the default pool, and then create an Exchange storage pool by adding only the arrays labeled 1 and 5 to the pool.

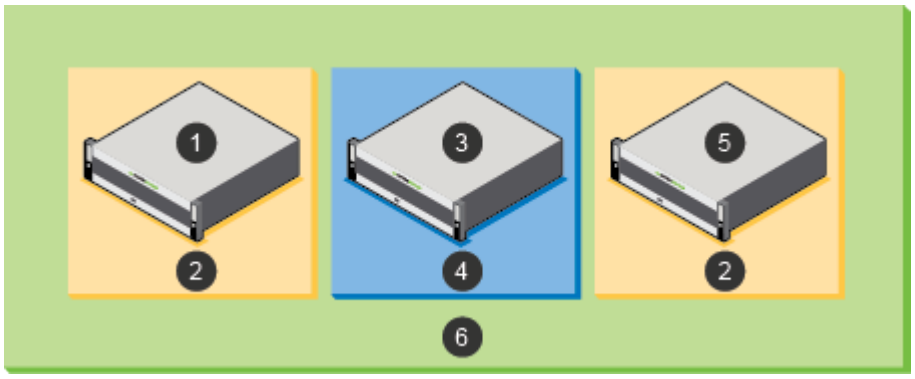

1 Array 1

2 Exchange pool

3 Array 3

4 Default pool

5 Array 5

6 Multi-array group that includes arrays 1, 3, and 5

The third implementation example shows adding a new array that is labeled 6 to the group that is assigned to the default storage pool. In this scenario, you have other alternatives. One alternative is to move array 6 to the previously created Exchange pool. Another alternative is to create a new storage pool for array 6 and leave only array 4 in the default pool.

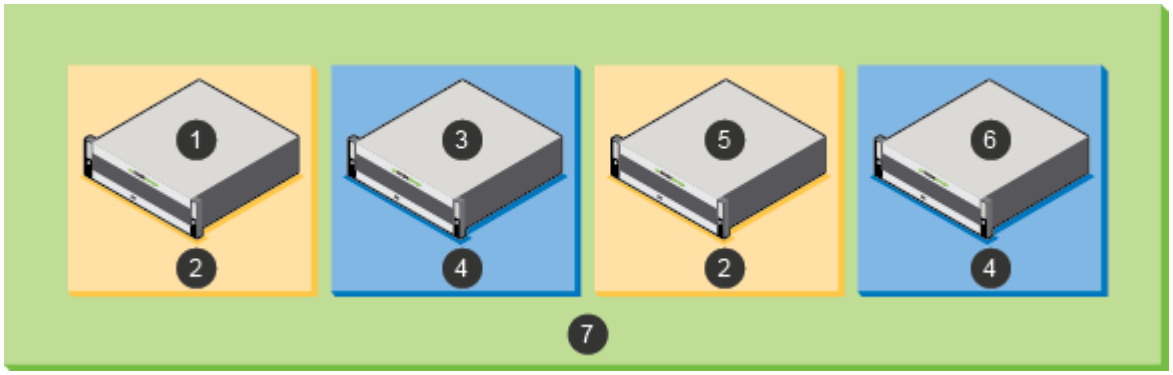

1 Array 1

2 Exchange pool

3 Array 3

4 Default pool

5 Array 5

6 Array 6

7 Multi-array group that includes arrays 1, 3, 5, and 6

The fourth implementation example shows having one group with four HPE arrays, where each array is in a separate storage pool.

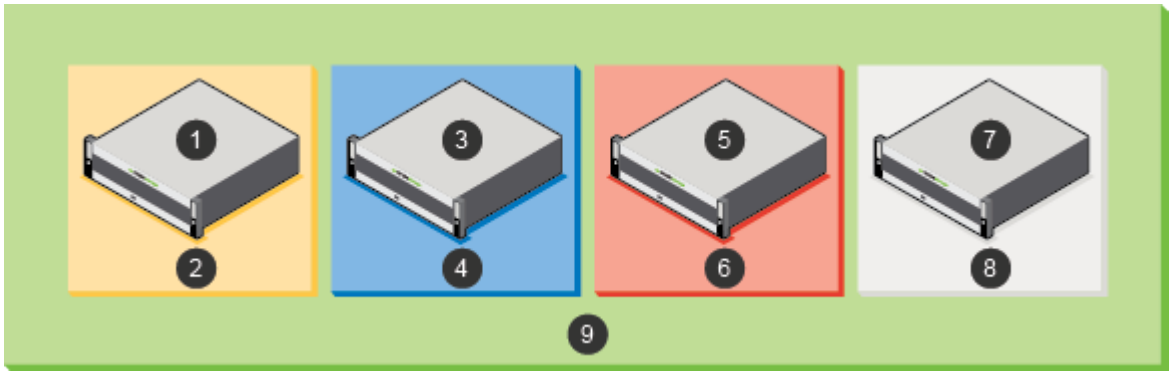

1 Array 1

2 Exchange pool

3 Array 3

4 Default pool

5 Array 5

6 SQL Server pool

7 Array 7

8 VDI pool

9 Multi-array group that includes arrays 1, 3, 5, and 7

Create a Storage Pool

By default, the arrays in a group can access each other and are grouped into a default pool. You cannot change the default storage pool, but you can add additional storage pools and move volumes between the pools.

A pool confines data to a subset of the arrays in a group. Data of resident volumes is striped and automatically rebalanced across all members of a group. Pools dictate physical locality and striping characteristics. Think of a pool as a logical container that includes one or more member arrays in which volumes reside. A member array and all of its expansion shelves can only be part of one pool.

Before you begin

You must have at least one unassigned array to put into the new storage pool.

You can also create a new storage pool when adding an array to a group.

Procedure

Create a storage pool.

pool --create *pool_name* **--description** *description* **--array** *array*

Example

Creating a pool with a description for the array c30-array6:

```
$ pool --create 1pool1 --description pool_for_admin --array c30-array6
```

Add or Remove Arrays from a Storage Pool

You can modify storage pools to add or remove arrays. Only one pool can be assigned to a given array. Multiple arrays can be assigned to the same pool.

Note: A pool must always have one member array.

CAUTION: Exercise caution when you change a storage pool assignment because it may result in a large migration of data.

Procedure

To add an array to a storage pool or remove an array from a storage pool, execute one of these commands:

- NimbleOS $ **pool --assign** *poolname* [**--array** *array*]

  or

- NimbleOS $ **pool --unassign** *poolname* [**--array** *array*]

Merge Storage Pools

You can merge two storage pools to combine the contents of both pools. To merge multiple pools, merge two pools at a time until the desired storage pool is created.

Procedure

Merge storage pools.

$ **pool --merge** *poolname* [**target** *target*]

Repeat this step for each of the pools you want to merge into the target.

Example

Example:

```
Nimble OS $ pool --merge Pool1 target Targetpool1
```

Volume Moves

Volume move operations allow you to move volumes between folders or between pools. When you move a volume from one pool to another, parts of the volume might exist on either pool. When the move is complete, the data exists only on the destination pool. It is not preserved on the source pool.

The two most common reasons to move data are to increase performance by moving a volume from a single-array pool to a multi-array pool and to manually balance space usage between pools.

When you move a volume, you also move all of its snapshots and clones. User access permissions and encryption settings move with the volume.

During a move, the change from one pool to another is transparent. There is no disruption of service, and during the move process, any write is made to the current location of that part of the volume. A move can require some time, depending on the amount of data in the volume, the number of clones, and the pools being migrated to and from.

Important: During a move, some volume data will reside partially on both pools; consequentially, when performing a move, you must use Nimble Connection Manager (NCM) on the host to facilitate the move. You should also use NCM to check hosts for old connections to targets that connect directly to data IPs. If such connections exist, then for each target:

- confirm that other connections to the corresponding group discovery IP exist, or
- establish new connections that point to the corresponding group discovery IPs before disconnecting these old connections.

In addition, make sure you perform the same checks for any defined static, favorite, or persistent targets and update any connections as necessary.

Volumes can be moved independently or simultaneously with other volume-move operations. However, while in the process of a move, you cannot start another move on the same volume.

> **Note:** For Fibre Channel groups, in order to avoid losing data paths between the host and the volume being moved, Nimble recommends that you check that Fibre Channel connectivity exists between the host and the arrays in a FC group prior to initiating a volume move operation.

Move a Volume from One Storage Pool to Another

Moving a volume also moves its snapshots and clones from one storage pool to another.

Depending on the size of the volume and the amount of data, it could take some time to complete a move from one storage pool to another. During this time, writes are allowed. Some forwarding of I/O requests from arrays in the source pool to arrays in the destination pool can occur during the move. This can result in temporarily reduced I/O throughput.

The destination pool must have sufficient space to accommodate current use and reserve space for all volumes being moved.

At any point during a move, an administrator can stop the move. When a volume move is stopped, all data that has already moved onto the destination pool moves back to the source pool. This may take some time, depending on the size of the pools and how much data has been moved.

Procedure

Move a volume to a destination pool.

vol --move *vollist* [**--dest_pool** *dest-pool*]

Example

Moving a volume to a new destination pool.

```
Nimble OS $ vol --move vol22 --dest_pool poolA
```

Move a Volume from One Folder to Another

Moving a volume also moves its snapshots and clones from one folder to another.

Depending on the size of the volume and the amount of data, it could take some time to complete a move. During this time, writes are allowed. Some forwarding of I/O requests from arrays in the source folder to arrays in the destination folder can occur during the move. This can result in temporarily reduced I/O throughput.

The destination folder must have sufficient space to accommodate current use and reserve space for all volumes being moved.

At any point during a move, an administrator can stop the move. When a volume move is stopped, all data that has already moved onto the destination folder moves back to the source folder. This may take some time, depending on the size of the folders and how much data has been moved.

Procedure

Move a volume to a pool and assign all volumes to a folder in the destination pool.

vol --move*volume_name* [**--dest_pool** *pool_name*] [**--dest_folder** *folder_name*]

Important:

See *vol* in the *Nimble Storage Command Reference* for more detail about this command.

Stop a Volume Move in Progress Using the CLI

You can stop an in-progress volume move at any point before completion. All data already moved into the destination pool moves back to the source pool.

Stopping a volume move can take some time, depending on the size of the pools and how much data has been moved.

If you need to undo a volume move after the move is completed, you must perform a second volume move back to the original pool.

Procedure

Stop the volume move.

vol --abort_move *name of volume*

Example

Example:

```
vol --abort_move vol22
```

Delete a Storage Pool

You can delete a storage pool that is obsolete or no longer used after you add its member arrays to other storage pools.

Note: The default storage pool cannot be deleted.

Before you begin

- Before you delete any other storage pool on an array, move its associated volumes to another pool or remove the volumes.
- An alternative to deleting a pool is to merge it into another pool on an different array. The result is that the source pool no longer exists and the destination pool contains any volumes that were in the source pool. There must be enough space in the destination pool for the space occupied by the volumes and snapshots of the source pool.

Procedure

Delete a storage pool.

pool --delete *pool_name* [**--force**]

The --force option forcibly deletes the pool even if it contains deleted volumes whose space is being reclaimed.

Example

Deleting a storage pool.

```
$ pool --delete MyPool
```

Data Protection

Ensuring that your data is protected is a critical part of managing Nimble arrays. Volumes that hold multiple components of an application, such as databases and transaction logs, can be grouped in *volume collections*. A volume collection includes a set of protection schedules to create snapshots of each associated volume at specified intervals.

The protection schedule specifies the downstream partner where the snapshots are sent. If a protection schedule is not used for replication, it does not have a downstream partner. For synchronous replication, the downstream partner specifies the location to which to synchronously replicate the data.

Snapshots are captured synchronously to assure a consistent backup of the entire application environment. Generally, each application uses a different volume collection.

The Nimble array ships with several predefined protection templates to use as the basis for volume collections. Protection templates include snapshot and replication schedules, as well as retention policies that are based on best practices for commonly used applications. You can copy the predefined templates and modify the copies, or create new protection templates to best match the requirements of applications used in your environment.

All volumes assigned to the volume collection use the same protection configuration of scheduled snapshots, replication, and retention settings. After the volume collection is created, it can be modified as needed.

Even if you plan to take snapshots of volumes manually, you can create a volume collection without schedules on those volumes. This capability is available in the command-line interface (CLI).

Volume Collections

Volume collections are sets of volumes that share data-protection characteristics, such as snapshot and replication schedules.

Snapshots for all volumes in a collection are captured synchronously to ensure that the data across these volumes is mutually consistent. Volumes in a collection share replication and snapshot schedules, as well as the retention policies for the snapshots. For disaster recovery, all volumes in a volume collection can simultaneously fail over to the replication partner.

Volume collections are limited to one downstream replication partner. If you need to replicate to multiple partners, create a separate volume collection for each replication partner you want to replicate to.

You can create a volume collection for each application. Volume collections can contain multiple schedules, which affect all volumes assigned to the collection. These schedules are additive, so when schedules overlap, snapshots and replicas are created for every schedule. Use synchronization when you create volume collection schedules that are specific to some Microsoft applications or VMware virtual machines to ensure consistency.

When configuring protection schedules, either as protection templates or as part of a volume collection directly, be sure that you allow enough time for one schedule to complete before another starts. For example, if Exchange protection schedule 1 has DB verification and replication enabled, and it runs every hour, and Exchange protection schedule 2 has DB verification turned off and no replication, and it runs every five minutes, it is possible that without enough time between snapshots in schedule 2, schedule 1 cannot start.

Careful planning can greatly enhance your backup strategy. You may want to retain hourly snapshots for a few days to allow quick recovery, and retain only daily snapshots for longer periods, such as several weeks. You can define retention periods for snapshots. The Nimble-provided protection schedules have default values that keep volumes in the normal retention range, which is 150 snapshots or fewer, retained either locally or on the replication partner. Although you can configure volumes with a high snapshot retention schedule, it is not recommended. As you define protection schedules, a meter identifies when you reach the high retention range. A mouse-over tooltip identifies the ramifications of defining protection schedules that exceed the normal retention range.

> **Note:** Volumes that have a high snapshot retention schedule count in the snapshot limits for the volume, group, and pool.

Create a Volume Collection

Create as many volume collections as needed.

> **Note:** If you have created a custom agent, specify the **generic** app type when creating a volume collection.

Procedure

1. Create a volume collection.

   volcoll --create *name* [**--prottmpl** *name*] [**--description** *text*] [**--app_sync** {*none* | *vss* | *vmware* | *generic*}] [**--app_server** *server*] [**--app_id** {*exchange* | *sql2005* | *sql2008* | *sql2012* | *exchange_dag* | *sql2014* | *sql2016* | *hyperv*}] [**--app_cluster_name** *cluster name*] [**--app_service_name** *service name*] [**--vcenter_hostname** *server*] [**--vcenter_username** *username*] [**--vcenter_password** *password*] [**--agent_hostname** *server*] [**--agent_username** *user_name*] [**--agent_password** *password*

   You can also create a volume collection from a protection template:

   volcoll--create *name* **--prottmpl** *name*

2. (Optional) Create a replication schedule for the volume collection.

   volcoll --addsched *name* **--schedule** *name*

3. (Optional) Create a snapshot for the volume collection.

   volcoll --snap *name* **--snapcoll_name** *name*

Example

Creating a volume collection:

```
Nimble OS $ volcoll --create VolColA
```

Creating a volume collection, giving it a description, and assigning a protection template:

```
Nimble OS $ volcoll --create VolColA --description CollectionA --prottmpl Pro▶
tectA
```

Creating a volume collection, specifying the application-specific synchronization required, naming the host server, and identifying the application running on the server:

```
Nimble OS $ volcoll --create VolColA --description CollectionA --app_server
Host1 --app_sync vss --app_id exchange
```

Creating a volume collection, and specifying a Windows application cluster and service name:

```
Nimble OS $ volcoll --create VolColA --description CollectionA --app_server
Host1 --app_id exchange app_cluster Cluster1
--app_service_name Mail
```

Creating a volume collection, and setting a VMware vCenter host name, user name, and password:

```
Nimble OS $ volcoll --create VolColA --description CollectionA --vcenter_hostname
 Host2 --vcenter_username James
--vcenter_password Password
```

Creating a volume collection replication schedule:

```
Nimble OS $ volcoll --addsched Replication1 --schedule Schedule1
```

Creating a volume collection snapshot:

```
Nimble OS $ volcoll --snap Snapshot1 --snapcoll_name Snapcoll1
```

Creating a volume collection with app sync type generic:

```
Nimble OS $ volcoll --create gen-sync-test-vc1 --app_sync generic --
agent_hostname 10.xx.xxx.xxx --agent_username Administrator --agent_password
Nim123#
```

Creating a volume collection with app id hyperv:

```
Nimble OS $ volcoll --create MyVolColl --app_sync
      vss --app_server 10.xx.xx.xxx --app_id hyperv
```

Protect a Standalone Volume

Some volumes may not need to be part of a volume collection. You can associate standalone volumes with volume collections to define their snapshot and replication schedules. The volume collection associated with a standalone volume cannot contain any other volumes.

There is no difference between a volume collection of a standalone volume and any other volume collection from the standpoint of the CLI.

> **Note:** You can remove a stand-alone volume collection or remove protection from a volume in a stand-alone volume collection by editing the volume and marking it as "Unprotected." Then the volume can be added to a volume collection that contains volumes that are not stand-alone.

Procedure

1. Create a volume collection.

   volcoll --create *name*

2. Associate the standalone volume with the volume collection.

   vol --assoc *name* [**--volcoll** *name*]

Example

Creating a volume collection:

```
$ volcoll --create Collection1
```

Associating a standalone volume with the volume collection:

```
$ vol --assoc VolumeA --volcoll Collection1
```

Modify a Volume Collection

You can modify a volume collection to make it more effective in a changing environment.

> **Note:** If you have created a custom agent, specify the **generic** app type when modifying a volume collection.

Procedure

1. Modify the volume collection.

   volcoll --edit *name* [**--newname** *name*] [**--description** *text*] [**--app_sync** {**none** | **vss** |**vmware** | *generic*}] [**--app_server** *server*] [**--app_id** {*exchange* | *sql2005* | *sql2008* | *sql2012* | *exchange_dag* | *sql2014* | *sql2016* | *sql2017* | *hyperv*}]

2. (Optional) Modify the replication schedule for the volume collection.

volcoll --editsched *name* **--schedule** *name* [**--newname** *name*] [**--repeat** *period*] [**--repeat_unit** {**minutes**|**hours**|**days**|**weeks**}] [**--at** *time*] [**--until** *time*] [**--days** *days*] [**--retain** *number*]

Example

Modifying the name and description of a volume collection, and the name of the replication schedule.

```
$ volcoll --edit Collection1 --newname Exchange1 --description ExhangeCollection1
```

Modifying a replication schedule and changing its name.

```
$ volcoll --editsched Protection1 --schedule Exchange1 --newname ExchangePro▶
tection
```

Editing volume collection while providing app sync type generic and its arguments:

```
$ volcoll --edit gen-sync-test-vc1 --agent_hostname 10.xx.xxx.xxx --
agent_username admin
      --agent_password Nim123
```

Delete a Volume Collection

You can delete any volume collection that has no volumes associated with it.

Procedure

1. Remove all volumes from the volume collection.

   vol --dissoc *name*

   Repeat this step until all volumes associated with the volume collection have been removed.

2. Delete the volume collection.

   volcoll --delete *name*

Example

Removing a volume from a volume collection, then deleting the volume collection.

```
$ vol --dissoc Volume1
$ volcoll --delete Exchange1
```

Protection Templates

Protection templates are sets of defined schedules and retention limits that you use to pre-fill the protection information when creating volume collections and standalone volumes. As a result, protection templates not only minimize repetitive entry of schedules, they also minimize errors and inconsistent setups by allowing the creation and management of a finite set of protection templates to meet all business needs.

After you create a volume collection, schedules and synchronization settings can be changed on the collection. This makes using the protection templates an easy, fast way to create multiple volume collections that share similar schedules: Use the same protection template to create as many volume collections as you want, then modify the volume collections with the changes specific to the needs of each collection. This means that you can create volume collections that are grouped as logical restoration groups.

You can create as many volume collections from the same protection template as desired. Later changes to the volume collection will not affect the template. Likewise, changes to the protection template do not affect previously created volume collections.

Note: You cannot edit or delete the predefined protection templates provided by Nimble Storage; however, you can create new protection templates as needed.

Important: When setting up protection schedule for a protection template, be sure that you allow sufficient time for one schedule to complete before another one starts. For example, if Exchange protection schedule 1 has DB verification and replication enabled and runs every hour, and Exchange protection schedule 2 has DB verification turned off and no replication and runs every five minutes, there might not be enough time between schedule 2 snapshots for schedule 1 to start.

Create a Protection Template

Create as many protection templates as needed. You can use the ones provided by HPE, or customize your own.

Note: If you create protection templates that use replication, configure both arrays as replication partners before creating the protection template.

Note: If you have created a custom agent, specify the **generic** app type when creating a protection template.

Procedure

1. Create a protection template, including a protection schedule.

   prottmpl --create *template_name* [**--description** text] [**--app_sync** {**none** | **vss** | **vmware** | **generic**}] [**--app_server** *server*] [**--app_id** {**exchange** | **sql2005** | **sql2008** | **sql2012** | **exchange_dag** | **sql2014** | **sql2016** | **sql2017** | **hyperv**}] [**--app_cluster_name** *cluster_name*] [**--app_service_name** *service_name*] [**--vcenter_hostname** *server*] [**--vcenter_username** *server*] [**--vcenter_password** *server*] [**--schedule** *name*] [**--repeat** *period*] [**--repeat_unit** {**minutes** | **hours** | **days** | **weeks**}] [**--at** *time*] [**--until** *time*] [**--days** *days*] [**--retain** *number*] [**--replicate_to** *partner*] [**--replicate_every** *number*] [**--num_retain_replica** *number*] [**--alert_threshold** *hh***:***mm*] [**--snap_verify** {**yes** | **no**}] [**--skip_db_consistency_check** {**yes** | **no**}] [**--disable_appsync** {**yes** | **no**}] [**--external_trigger** {**yes** | **no**}]

2. (Optional) Add another protection schedule to the specified protection template.

   prottmpl --addsched *template_name* **--schedule** *name*] [**--repeat** *period*] [**--repeat_unit** {**minutes** | **hours** | **days** | **weeks**}] [**--at** *time*] [**--until** *time*] [**--days** *days*] [**--retain** *number*] [**--replicate_to** *partner*] [**--replicate_every** *number*] [**--num_retain_replica** *number*] [**--alert_threshold** *hh***:***mm*] [**--snap_verify** {**yes** | **no**}] [**--skip_db_consistency_check** {**yes** | **no**}] [**--disable_appsync** {**yes** | **no**}] [**--external_trigger** {**yes** | **no**}]

   Note: You can create up to ten protection schedules for the specified protection template. For example, you can create one schedule for working hours, one for peak hours, and one for weekends. Schedules can overlap, but they cannot span midnight.

Example

Creating a new protection template named highrepl.

```
prottmpl --create highrepl --description "use when high replication is needed"
--app_sync none --schedule replicated --repeat 1 --repeat_unit day
--days Monday,Tuesday,Thursday --retain 8 --replicate_to greyhound
--replicate_every 1 --num_retain_replica 40 --snap_verify no
```

Adding an hourly schedule to the highrepl protection template.

```
$ prottmpl --addsched highrepl --schedule hourly --repeat 1 --repeat_unit hours

    --retain 25 --snap_verify no
```

Creating a protection template with app sync type generic.

```
$ prottmpl --create gen-sync-test-prottmpl --app_sync generic --agent_hostname
 10.18.237.121 --agent_username Administrator --agent_password Nim123# --
schedule sched1 --retain 30
```

Edit a Protection Template

You can edit protection templates, but the changes will apply only to volume collections based on the template after it has been changed. Existing volume collections are not affected.

> **Note:** If you have created a custom agent, specify the **generic** app type when editing a protection template.

Procedure

1. Edit a protection template and at least one parameter to be changed.

   prottmpl --edit *template_name* [**--name** *new_name*] [**--description** text] [**--app_sync** {**none** | **vss** | **vmware** | **generic**}] [**--app_server** *server*] [**--app_id** {**exchange** | **sql2005** | **sql2008** | **sql2012** | **exchange_dag** | **sql2014** | **sql2016** | **sql2017** | **hyperv**}] [**--app_cluster_name** *cluster_name*] [**--app_service_name** *service_name*] [**--vcenter_hostname** *server*] [**--vcenter_username** *server*] [**--vcenter_password** *server*]

2. (Optional) Edit an existing schedule assigned to the protection template.

   prottmpl --editsched *template_name* **--schedule** *schedule_name* [**--repeat** *period*] [**--repeat_unit** {**minutes** | **hours** | **days** | **weeks**}] [**--at** *time*] [**--until** *time*] [**--days** *days*] [**--retain** *number*] [**--replicate_to** *partner*] [**--replicate_every** *number*] [**--num_retain_replica** *number*] [**--alert_threshold** *hh*:*mm*] [**--snap_verify** {**yes** | **no**}] [**--skip_db_consistency_check** {**yes** | **no**}] [**--disable_appsync** {**yes** | **no**}] [**--external_trigger** {**yes** | **no**}]

Example

Editing the protection template named highrepl to replicate once a week, and to retain 31 snapshots and 52 replicas.

```
$ prottmpl --edit highrepl --schedule replicated
--repeat 1 --repeat_unit weeks --retain 31 --replicate_to array7
--replicate_every 1 --num_retain_replica 52 --snap_verify no
```

Editing a protection template with app sync type generic.

```
prottmpl --edit gen-sync-test-prottmpl --agent_hostname 10.xx.xx.xxx
```

Delete a Protection Template

Over time, you may find that a template is no longer needed.You can delete any user-defined protection template without affecting existing volume collections.

Procedure

Delete a specified protection template.

prottmpl --delete *template_name*

Example

Deleting the protection templae named highrepl.

```
$ prottmpl --delete highrepl
```

Snapshots

You can manage snapshots the same way that you manage volumes. In reality, snapshots are volumes. They are subject to the same controls and restrictions as volumes. You can clone, replicate, and modify snapshots. Initiators can access snapshots.

Snapshots Overview

The initial snapshot for a volume uses no space because it shares its original data with the volume from which it was taken. Each successive snapshot consumes some amount of space because it captures the changes that occurred on the volume. The changed blocks are compressed to reduce capacity consumption.

Consider the change rate of the applications using the volume, the assigned replication strategy, and the snapshot retention to determine the amount of space you need for snapshots. You can retain numerous snapshots in most environments.

A *snapshot reserve* represents the amount of space that you allocate (pre-reserve) from the general pool for snapshots of a volume. Each successive snapshot captures and stores changes to the volume contents. Snapshot reserve space can be increased if needed. When the snapshots exceed the amount of space reserved, space is then used from the general storage pool.

The optional *snapshot quotas* determine how much of the volume can be consumed by snapshots. For example, you may want to allow no more than 10% of the volume space to be used for backup. Setting the snapshot quota to 10% ensures this volume is not exceeded. The actual amount of space you need depends on the application using the volume and the volume collection to which it is assigned.

If a volume is approaching the quota limit, an event is logged. When enforcement is enabled, the volume administrator can determine what follow-up actions to take, such as preventing users from accessing more disk space or allocating additional disk space.

Initiators access online snapshots just like they access online volumes. To access snapshot data, set the snapshot online and log the initiator into the snapshot.

When you delete a volume, the snapshots that are associated with that volume are also deleted. If the volume has online snapshots, they must be taken offline before you can delete them.

Snapshots and Daylight Savings Time

In cases where the system time changes, whether it is changed manually or automatically, if the time change is identical to a protection schedule interval, the next scheduled snapshot is skipped. Daily schedules may have two snapshots, and are skipped only if the snapshot time falls within the missed time frame. For example, if you have an hourly snapshot schedule and the system makes a Daylight Savings Time adjustment of one hour, the next scheduled snapshot is skipped because the system sees the appropriate snapshot for that hour is already taken.

Snapshot Rate Limits

All volumes in a group are associated with a volume collection. Multiple volume collections may be created for each group, and each volume collection may have multiple protection schedules. Protection schedules may have different periods, or rates, such as hourly, daily, or weekly. An hourly schedule takes a snapshot of all volumes in the volume collection once per hour. A schedule that is configured to start at midnight begins taking snapshots of all volumes in the collection at 00:00, and at each hourly interval thereafter.

Beginning with OS 4.x, the maximum expected snapshot completion rate is 250 snapshots per minute per array group. The maximum cumulative outstanding snapshot count is the total of all snapshots for all protection schedules in the group, and cannot exceed 4,000 at any point in time. The outstanding snapshot count for a schedule is the number of snapshots started

at that point in time or during the prior minute, but not completed because of the 250 snapshot per minute rate limitation. Outstanding snapshots are carried forward to the next minute, and completed at a rate of 250 per minute until all are finished.

For example, a protection schedule is configured to start and repeat every five minutes (T equals 5), and contains 270 volumes. The outstanding snapshot count at minute one is 270. All snapshots in the schedule must be completed in five divided by two (T/2), or three minutes. During minute one, 250 snapshots are taken. The remaining 20 snapshots are carried forward and completed during minute two. When the schedule repeats at the next five-minute interval, 250 snapshots are taken during minute five and the remaining 20 during minute six.

Schedules that have a lower repeat interval, which is calculated as the number of minutes in the interval, have a higher priority than schedules with a higher repeat interval. For example, daily schedules are prioritized to start before hourly schedules. When multiple schedules have the same repeat interval and are configured to start at the same time, the one that was created first has the higher priority. Higher priority schedules are considered first in determining the maximum outstanding snapshot count and the total expected snapshot completion time.

The more aggressive the schedule, the fewer volumes that can be protected. In order to prioritize all snapshots, those with daily schedules are completed before those with hourly intervals. Snapshots with hourly schedules are completed before those with intervals in minutes.

If you create ten volume collections, each with a single hourly protection schedule containing 50 volumes, a total 500 snapshots are scheduled. An hourly schedule has a time (T) period of 60 minutes. All snapshots must be completed in T divided by two, or 30 minutes. If all ten schedules are set to start at the same time, a maximum of 250 snapshots are taken during the first minute of the hour, and 250 are carried forward. The remaining 250 snapshots are taken during the second minute.

If you then create ten new volume collections, each with a daily protection schedule containing 50 volumes, the 500 new daily snapshots are added to the 500 hourly snapshots. You now have a cumulative total of 1,000 snapshots. If all ten daily schedules are scheduled to start at the same time as the ten hourly schedules, the snapshots are taken in the following order:

1. The first 250 of the daily snapshots are complete in minute 1, and the remaining 250 daily plus the 500 hourly are carried forward.
2. The remaining 250 daily snapshots are completed in minute 2, and the 500 hourly snapshots are carried forward.
3. The first 250 of the hourly snapshots are completed in minute 3, and the remaining 250 hourly snapshots are carried forward.
4. The final 250 hourly snapshots are completed in minute 4.

When you create or edit snapshot schedules, the OS makes calculations to determine whether the changes will exceed the 250 snapshots per minute or the maximum outstanding snapshot count 4,000 limits. If either one will be exceeded, the operation fails. To avoid exceeding the outstanding snapshot count when you add a new schedule or edit an existing schedule, stagger the schedule start time.

Volume and Snapshot Usage

Space used by a snapshot is never more than the space used by the volume at the time the snapshot was created. Any "live" new data introduced by snapshots is attributed to the live volume space usage and not the snapshot usage. Blocks are attributed to snapshot space usage when they are overwritten in live state. The easiest way to think about this is that snapshots consume space only when a block is modified or deleted.

For example, if you write 50 GB to a volume and take a snapshot, the initial snapshot shows zero snapshot usage, because it points to the same blocks on disk, which are "live." If you write another 50 GB to the volume, making a total of 100 GB, and then take another snapshot, the snapshots still consume no space. In fact, you can create several snapshots and they will all consume no additional space for the newly written 100 GB.

If you overwrite the 100 GB, then all snapshots start to consume space. Their cumulative space usage should be approximately 100 GB.

When some applications, such as PowerPoint, Excel, and Word modify files, the applications do not necessarily update the specific blocks in the file that changed. Instead, the applications create a new file. In this case, after you modify a file from one of these applications, a snapshot could show no usage because the original file the snapshot points to was not changed, but instead was seen as a new file.

Automatic and Manual Snapshots

You can manually take a volume snapshot at any time. Manual snapshots are often used for testing a new application before integrating it with a production volume or for troubleshooting.

You can assign a volume to a volume collection so that snapshots and replication automatically occur based on the schedules associated with the volume collection. Automatic snapshots are typically used for backup operations. Use your existing backup software, triggered by the snapshot schedule.

> **Note:** When the system time changes, and the time change is identical to a protection schedule interval, the next scheduled snapshot is skipped. Daily schedules may have one or more snapshots per day, and are skipped only if the snapshot time falls within the missed time frame. For example, if you have an hourly snapshot schedule and the system makes a Daylight Savings Time adjustment of one hour,the next scheduled snapshot is skipped because the system identifies that the snapshot for that hour was already taken.

Take a Manual Snapshot

You may need to take an on-demand (manual) snapshot of a volume before you update software or make hardware changes on the array.

Even if you plan to take snapshots of volumes manually or use a third-party program, you can create a volume collection without schedules on those volumes for which snapshots are being manually taken. You may see a data service restart abort message if you take a snapshot around the same time that a scheduled snapshot for the volume is going to start. Check the volume collection to ensure that no snapshot is pending before you trigger a manual snapshot.

> **Note:** Manual snapshots are not guaranteed to be application consistent.

Procedure

Take a manual snapshot.

vol --snap *vol_name* [**--snapname** *name*] [**--description** *text*] [**--start_online**] [**--allow_writes**]

Example

Taking a snapshot of a volume:

```
$ vol --snap Vol1 --snapname Snap1
```

Taking a snapshot of a volume with a description, and setting access to allow applications to write to the snapshot:

```
$ vol --snap Vol1 --snapname Snap1 --description Daily --allow_writes
```

Clone a Snapshot

Clones of are useful for restoring individual files instead of a complete volume. By copying the files from the cloned volume to the active volume, you can restore the files without affecting other users.

Procedure

Clone volume data from a specific snapshot.

vol --clone *vol_name***--snapname** *snap_name* **--clonename** *clone_name* [**--description** *text*]

A clone of the snapshot is created.

Example

Cloning volume Test1 data from snapshot Snap1 to the clone volume CloneRestore:

```
$ vol --clone Test1 --snapname Snap31 --clonename CloneRestore
```

Change a Snapshot's State

When snapshots are created, they are set offline. Setting a snapshot offline makes it unavailable to initiators, and closes any current connections.

Note: Snapshots that are not required to be online should be kept offline until needed.

Procedure

Change the state of a snapshot.

Use one of the following commands, depending on the current state of the snapshot:

Option	Description
snap --online *snap_name*--vol *vol_name*	Brings the snapshot online.
snap --offline *snap_name*--vol *vol_name* [--force]	Takes the snapshot offline. The [--force] option forcibly disables access to the snapshot.

Delete a Snapshot

Unlike deleting a volume, deleting a snapshot has no impact on the original volume. Only the data on the snapshot is lost.

You cannot delete a snapshot while it is online. Set the snapshot to offline before you delete it.

Procedure

1. Take a volume snapshot offline.
   snap --offline *snap_name* **--vol** *vol_name* [**--force**]

   The --force option forcibly disables access to the snapshot.
2. Delete the snapshot.
   snap --delete *snap_name* **--vol** *vol_name* [**--force**]

   The --force option forcibly deletes a snapshot that is managed by an external agent.

Hidden Snapshots

When you install an array, you typically set up volumes, as well as volume collections and snapshot schedules. If you do not set up any volume collections or snapshots, the array automatically generates a snapshot every hour. These snapshots are termed hidden snapshots.

While individual hidden snapshots are not listed under the Snapshot tab, hidden snapshots usage is part of the calculation of the Snapshot Usage column on the Space tab. Similarly, if you schedule snapshots that occur more than one hour apart, the array continues to generate hidden snapshots. As soon as you decrease the frequency of snapshots to something less than one hour, the array stops taking hidden snapshots.

Snapshot Consistency

Stagger snapshot schedules to ensure application synchronization, I/O quiescing, database verification, and so on. Consider the following points for different application types.

Microsoft application snapshots	For some Microsoft applications, such as Microsoft Exchange®, snapshots require that the application writes are flushed to the database and traffic is stopped while the snapshot is taken. This ensures that there is never partial data stored in the snapshot. The OS performs this step automatically when Microsoft VSS synchronization is enabled.
VMware snapshots	If your data center uses VMware vCenter, ensure that traffic is stopped while the snapshot is taken so that the snapshot is complete and can be cloned directly to a new virtual machine. The OS performs this step automatically when VMware vCenter synchronization is enabled.
Application-consistent snapshots with VMFS	The VMware snapshot captures the state and data of a virtual machine at a particular point in time. When creating a snapshot, VMware provides the "quiesce" option which flushes dirty buffers from the guest OS in-memory cache to disk, and offers application consistency through VSS requestor in VMware Tools. The Protection Manager takes advantage of the VMware quiesced snapshot option and combines it to achieve consistent and usable volume snapshots and replicas.

Snapshot Framework

The Snapshot Framework allows you to write custom host- or application-aware plug-ins (also known as “agents”) to customize the pre-snapshot and post-snapshot tasks. By default, the array provides application-consistent snapshots and replication of vSphere datastores, MS-Exchange, MS-SQL, and NTFS on the following platforms:

- VMware (through vCenter synchronization)
- Microsoft SQL Server (through Microsoft VSS sync)
- Microsoft Exchange Server (through Microsoft VSS sync)

However, for applications that are not VSS-aware, a custom plug-in created with the Snapshot Framework can be used.

The Snapshot Framework dramatically expands the set of applications that can be integrated with HPE Storage Snapshots, including Linux Oracle and SAP applications, and even Windows applications that are not VSS-aware.

The Snapshot Framework does not replace VSS Integration. Any third-party backup applications that are VSS-aware can integrate with the OS normally via the Storage VSS provider; non-VSS backup applications can use REST APIs.

For information on how to develop your own agent, refer to the *Snapshot Framework Reference*.

You can use the array OS CLI to perform the following tasks using your custom agent:

- Create a Volume Collection on page 92
- Modify a Volume Collection on page 93
- Create a Protection Template on page 95
- Edit a Protection Template on page 96

NSs Snapshots

NSs snapshots are temporary snapshots that have NSs-* prepended to the snapshot name. These snapshots are commonly created in the following scenarios:

- Volume-level restores
- Volume size increase
- Volume size decrease
- VSS related snapshot operations

Normally, NSs snapshots are set to offline; however, the snapshots might appear online if the process using them did not complete successfully.

In rare cases, and more commonly in legacy code, if an unexpected failure occurs that prevents NSs snapshots from being cleaned up, you might see an NSs snapshot left online even though there are no running back end processes that leverage the snapshot.

NSs snapshots are usually renamed upon successful backup to meet the normal snapshot-naming scheme; however, if a third-party backup requester calls our provider, this could cause the snapshot to retain the NSs-* name.

> **Note:** NSs snapshots taken with array-side volume activities are considered unmanaged snapshots; once the snapshots are no longer needed, they can be cleaned up manually.

Working with Online Snapshots

Online snapshots are useful for verifying the contents of a snapshot. They can be generated either automatically or manually.

Note: It is recommended that you use VSS verified snapshots instead of manually creating an online, writable snapshot. Depending on the type of data on the volume, having VSS enabled is necessary to maintain recoverable data.

Certain backup utilities and data scanning utilities automatically create online snapshots. Normally these snapshots remain online only for the duration of the operation. If utility does not turn off the online feature after the process finishes, you should manually turn it off.

You should **not** use online, writable snapshots as a replacement for cloning. For example, you should not use online snapshots in the following situations:

- Do not use online writable snapshots to add storage to a host by making it an extension of the existing file system.
- If you plan to use a snapshot for testing or some situation where you need to use the snapshot as a regular volume. Instead, you should create a clone and use that a new volume and migrate data as needed.

Online, writable snapshots do not maintain point-in-time (PIT) information. If you change the data on an online snapshot, those changes will persist when the snapshot is turned offline. The changes to the upstream online snapshot will not replicate and the replica snapshot will only contain the data that was used to the create the original online snapshot.

If you create an online, unwritable snapshot, the data in the snapshot cannot be changed, so you do have a PIT snapshot.

> **Note:** If you manually create an online snapshot and you make it non-writable, you cannot change it later to make it writable.

Identify Online Snapshots Using the NimbleOS CLI

You can identify snapshots that are online by using the NimbleOS CLI.

Procedure

1. Enter the command **snap --list --vol <***volume_name***> | awk '$4 == "Yes"'**

   This command displays only the online snapshots for the volume you specify with *volume_name* parameter .

2. If you are running NimbleOS 2.3.x or later, you can narrow the snapshot list down so that it displays only the online unmanaged snapshots on a specific volume. Enter:

   snap --list --unmanaged --vol <volume_name**> | awk '$4 == "Yes"'**

3. Determine the state of the snapshot and whether it has actively connected initiators by entering:

   snap --info <*snapshot_name***> --vol <***volume_name***>**

4. 4. If there are active connectors, complete the operation that is in progress; for example migrating or moving data.
5. Determine whether there is a reason to have the online snapshot in the future:

   - Is the snapshot required by an application or script. If it is, take the following actions:

     1 Check the retention schedule to make sure the snapshot will be removed at a future date.
     2 Disconnect the snapshot from the host before the next snapshot deletion operation is scheduled to occur.

   - Is this the only common snapshot between the upstream array and the downstream array. If it is, then deleting the snapshot would cause the volume to need a full re-seed.

6. If you do not need to maintain an online snapshot, perform the following steps:

   1 Disconnect the snapshot from the host.
   2 Make it offline.
   3 Delete it.

Migrate Data From an Online Snapshot to a New Volume

If you are using an online, writable snapshot as an extent to an existing file system, it is recommended that you create a new volume and migrate the data from the snapshot to it.

Procedure

1. Create a new volume.
2. Migrate the data from the online snapshot to the new volume.
3. Connect the host to new volume.
4. Verify that all data migrated successfully.
5. Gracefully disconnect the initiators from online snapshot.
6. Turn the snapshot offline and, if it is no longer needed, delete it.

Replication

You can use volume replication to copy critical data to Nimble arrays at different locations as part of your disaster recovery strategy. Replication does not take the place of snapshot backups, but it enhances the overall data recovery plan.

For data recovery tasks like recovering an accidentally deleted or corrupted file, you can take snapshots that serve as a backup. Snapshots have little performance impact, can be performed quickly, and are space efficient.

For more widespread issues like a power failure or a site disaster, replication technology can be a quick and effective way to recover data at an offsite location. In these scenarios, data can be served from the replica while the initial array is being restored.

Nimble arrays use an advanced file system that provides in-line compression for data writes. Data is stored in variable-length blocks that match the logical write methodology of an application. These features minimize the amount of replication traffic to a compressed version of just the logical application write and reduce bandwidth requirements.

Replication tasks are scheduled automatically through protection templates. You can implement any number of replication strategies to meet your requirements for disaster recovery. For more information on protection templates, see the chapter on *Data Protection*.

Replication Overview

You can use volume replication to copy critical data to Nimble arrays at different locations as part of your disaster recovery strategy. Replication does not take the place of snapshot backups, but it enhances the overall data recovery plan.

What is Replication?

Replication maintains a copy or replica of a volume or set of volumes and their snapshots on another Nimble array that is configured as one of a pair of replication partners. The replica contains the contents of a volume at the time the replica was created or last updated, as well as a configured number of prior states (snapshots). Nimble replication is based on snapshots. Replicas are stored at a remote array, called a replication partner, connected by a network or Internet link. A volume is always located on a different array than its replica. It is possible to retain more or fewer snapshots on the replica than on the source volume, thus providing greater flexibility in designing a recovery plan.

A replica is a copy of a volume from another group whose state is managed by the other group. For example, all write operations to a replica originate from another group. The group that hosts the original volume and manages the state of a replica is called the upstream partner, because the data flows from it. The group where the replica resides is called the downstream partner, because the data flows to it.

The replicated volumes can be restored as complete copies of the volumes, with all schedules and administrative settings replicated, as well as the actual data. Like snapshots, replicas are created and stored based on volume collection schedules. Multiple volumes can share a volume collection schedule.

When you promote a replica, the number of snapshots to retain is adjusted to be the maximum number of snapshots to retain on the local array, plus the number of snapshots to retain on the replica. The volume is offline until a replica is promoted. The volume and settings are visible on the replica, but they are not editable until a volume collection handover is performed.

Replication Partners and How Replication Works

You can set up replication between two Nimble arrays or a Nimble array and HPE Cloud Volumes (HPE CV). By creating a partner on each of the two arrays, replication is automatically set up. You can configure Nimble arrays to have up to fifty replication partners (also known as replication stores in HPE CV). Replication is automatic, based on the protection schedule assigned to the volume or volume collection.

Replication partners can be reciprocal, upstream (the source of replicas), or downstream (the receiver of replicas) partners. You can have several upstream arrays that replicate to one downstream partner.

When you create a replication partner, you enable one Nimble array to perform replication tasks on another array. The two arrays must be able to communicate over a network.

Note:

- Replication uses ports 4213 and 4214. Replication of encrypted volumes uses port 5391. See Configure Firewall Ports on page 230 for more information.
- Replicating encrypted volumes to HPE CV requires NimbleOS 5.0.6.x or higher.
- Replication to HPE CV is supported only on iSCSI arrays. Replication with Nimble arrays can use different access protocols. One of the partners can be an iSCSI array and its replication partner can be a Fibre Channel array.

Replication partners can use different access protocols. One of the partners can be an iSCSI array and its replication partner can be a Fibre Channel array.

Note: If you create protection templates or volume collections that use replication, configure both arrays as replication partners before you begin.

Volumes are replicated by associating them with volume collections with at least one replication schedule. A volume collection can have up to ten protection schedules. One or more of these protection schedules can be replicating schedules. All replicating schedules of a volume collection are replicated to the same replication partner.

Using volume collection schedules, you define which replication partner and at what times volumes should be replicated, and how many point-in-time versions of the replica should be retained. You can assign volumes to the necessary volume collections. At the scheduled time, the array from which the volume is being replicated sends the data to the downstream (receiving) replication partner. You can also create a new protected standalone volume that is automatically associated with a volume collection instead of creating a volume first and then associating it with a volume collection.

If a volume uses a custom performance policy, you must duplicate that policy on the replication partner.

When you enable replication of a volume for the first time, the entire data set on the volume at the time of the snapshot is replicated from the original array group (upstream) to the destination replication partner (the downstream array group). Subsequent updates replicate the differences between the last replicated snapshot on the upstream partner to the downstream partner.

It is recommended that you configure the snapshot retention count on the replica (downstream) array group to at least the number of snapshots on the upstream array group, matching the frequency of snapshots.

Note: In majority of NimbleOS releases, by default, the replication partner retains only two snapshots for each replica volume in the volume collection if configuration is not adjusted.

Snapshot collections are replicated in the order that the collections were taken. Once replication is caught up, the upstream and downstream replication partners only retain as many snapshots as set in the retention criteria. The system deletes pending snapshot collections that exceed the retention criteria.

Create a Replication Partner

Use the management subnet for replication when any of the following conditions apply:

- Your data IP addresses are not routable across the network.
- You want to separate replication traffic from iSCSI traffic.
- Your replication HPE Nimble Storage arrays are running NimbleOS 1.4 or NimbleOS 2.x or above.

Note: For NimbleOS 2.x and later, replication over a data subnet is available; however, it requires the replication control traffic to be transferred over a management subnet. If you choose to replicate over a data subnet, you must be able to do the following:

- Route the management subnet between replication partners by the default gateway for replication control traffic.
- Route the management subnet between replication partners by the default gateway for replication control traffic. Route the data subnet between replication partners by a static route for replication data transfer traffic.

Replication partners can run different data access protocols. For example, a Fibre Channel (FC) replication partner can be created for an iSCSI array, and an iSCSI replication partner can be created for a FC array.

Note: For FC arrays, you do not need to separate the replication traffic; replication traffic never runs over FC.

Procedure

1. Determine the IP address or hostname of the array to be used as a replication partner.
2. Add a replication partner.

   partner --create *partner_name* **--hostname** [*ipaddr*|*hostname*]**--description** *text* **--secret** *shared_secret*

   - For *partner_name*, you must enter the current Group name of the partner. This value is case-sensitive.
   - The description is optional.
   - The shared secret must be eight or more characters with no spaces or special characters. Special characters include: ' " ` ~ ! @ # $ ^ & ( )+ [ ] {} * ; : ' " . , | < > ? / \ = %.

3. (Optional) Select the replication network.

   partner --edit *partner_name***--subnet** *local_subnet_label*

   Use the --subnet option only when there are multiple data subnets. When you use the *local_subnet_label* variable, include the full subnet address and mask, separated by a slash. For example, 172.18.120.0/255.255.255.0

4. (Optional) Create a QoS (replication traffic bandwidth limit) policy.

   Note: A Quality of Service (QoS) policy defines how the network resources are allocated for the replication partner. Without a QoS policy, the replication partner can use unlimited network bandwidth.

   partner --create_throttle*partner_name* **--description** *text* **--days** *days* **--at** *start_time* **--until** *end_time* **--bandwidth** *limit*

What to do next

You must log in to the replication partner and perform the same configuration on that array and group.

Modify a Replication Partner

You can modify the IP address or bandwidth requirements for a replication partner without having to recreate the replication partner. You cannot change the name of a replication partner. In this case, create a new replication partner with the required name, assign the volume collection to the new replication partner, and delete the original replication partner.

Procedure

Modify the replication partner.

partner--edit

partner--edit_throttle

Delete a Replication Partner

Note: Before you delete a replication partner, you must perform the following actions:

- Change all volume collections so that replication is not scheduled with a partner
- Ensure that there are not any volume collections that have the replication partner selected

  You can determine this by searching by the partner name on the volume collections page.

- Be sure to remove replication partners on both the upstream and downstream partners

When you delete a replication partner relationship, only the replication partner relationship is deleted, not the array.

Procedure

Delete a replication partner.

partner --delete *partner_name*

If you removed all volume collection schedules from the source group, you can view the confirmation that the replication partner was deleted. If there are still active schedules associated with the partner, delete those and repeat the process. You may want to promote the volume collection on the downstream partner before you delete it.

Test the Connection between Replication Partners

You can test the connection between configured replication partners at any time.

Procedure

Test the connection between replication partners.

partner --test *partner_name*

Replication Strategy

Several options for replication strategy are available. Each has advantages and disadvantages. You need to decide on the best strategy for your environment. For example, you might use different configuration options that are based on available space, application, criticality of data, and legal requirements. Consider your environment, applications, availability needs, storage growth patterns, and recovery windows to create a replication strategy that best serves your needs.

The Nimble array relationships are based on the direction of replication. The source of the volume collection that is being replicated is the *upstream* replication partner. The destination is the *downstream* replication partner. A volume source that is upstream to one partner can be downstream to another. Downstream partners can become upstream partners by sending data to other replication partners and by replicating replicas.

One-to-One Replication

Basic replication involves replicating volumes from one array to another based on the protection schedules configured for their associated volume collections. Each volume collection always replicates to the same replication partner. In this scenario, the second Nimble array could be used strictly for disaster recovery.

In this example, Hourly and Daily represent volume collections that are configured with protection schedules that take snapshots at the specified frequency. Volume collections can be configured with multiple protection schedules.

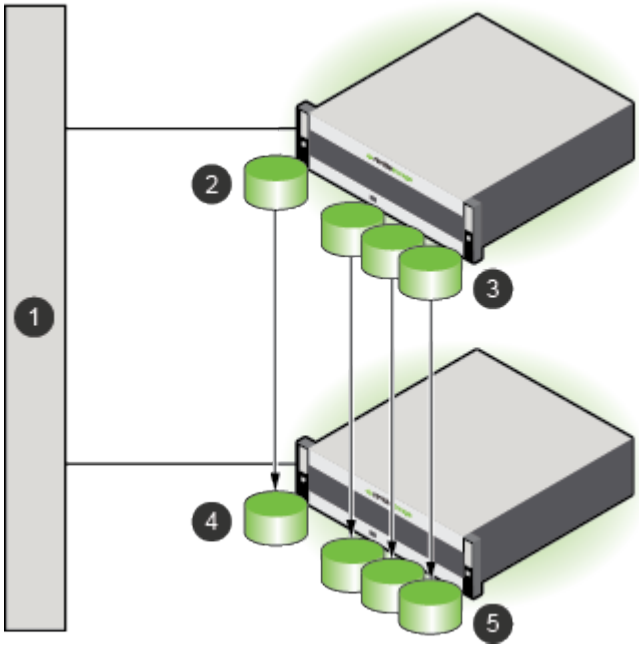

1 Network

2 Single volume assigned to the Hourly volume collection

3 Multiple volumes assigned to the Daily volume collection

4 Replica of volume in the Hourly volume collection

5 Replicas of volumes in the Daily volume collection

Reciprocal Replication

Reciprocal replication involves replicating volumes that originate on two separate arrays to each other. Volumes on one Nimble array are replicated to a second Nimble array, and volumes on the second Nimble array are replicated to the first array. Each array acts as a disaster recovery option for the other. Reciprocal replication is sometimes called *mutual replication*.

In this example, SQL and Outlook represent volume collections that are configured with protection schedules and performance policies appropriate for those application types.

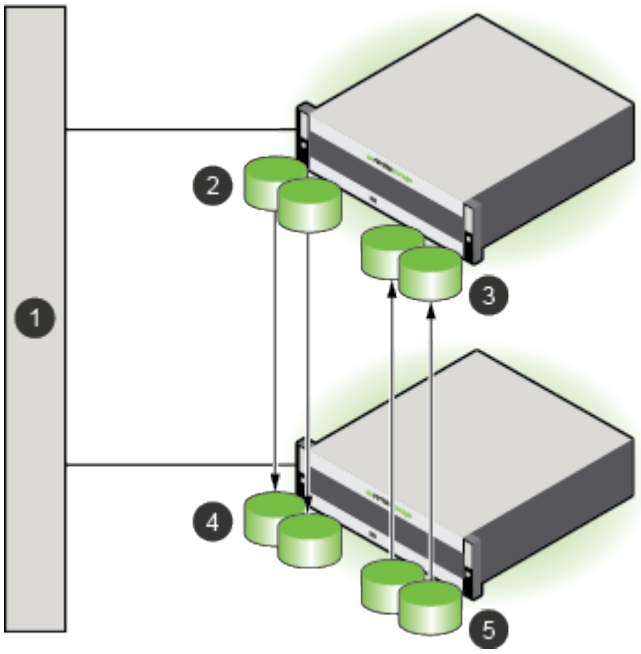

1 Network

2 Multiple volumes assigned to the SQL volume collection

3 Replicas of volumes in the Outlook volume collection

4 Replicas of volumes in the SQL volume collection

5 Multiple volumes assigned to the Outlook volume collection

Many-to-One (Centralized) Replication

You can use one Nimble array as a centralized replica for volumes that originate on several other Nimble arrays.

In this example, Hourly, Daily, SQL, Outlook, and Datastore1 represent appropriately configured volume collections.

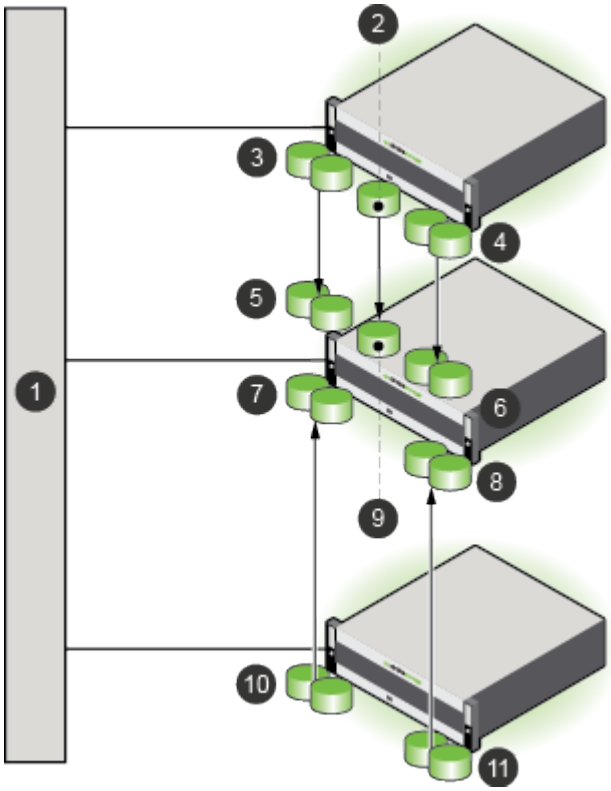

1 Network

2 Single volume assigned to the Daily volume collection

3 Multiple volumes assigned to the SQL volume collection

4 Multiple volumes assigned to the Hourly volume collection

5 Replicas of volumes in the SQL volume collection

6 Replicas of volumes in the Hourly volume collection

7 Replicas of volumes in the Outlook volume collection

8 Replicas of volumes in the Datastore1 volume collection

9 Replica of volume in the Daily volume collection

10 Multiple volumes assigned to the Outlook volume collection

11 Multiple volumes assigned to the Datastore1 volume collection

Many-to-Many Replication

Many-to-many replication involves replicating volumes on the same array to multiple replication partners.

In this example, Hourly, Daily, SQL, Outlook, Datastore1, and Temporary represent appropriately configured volume collections.

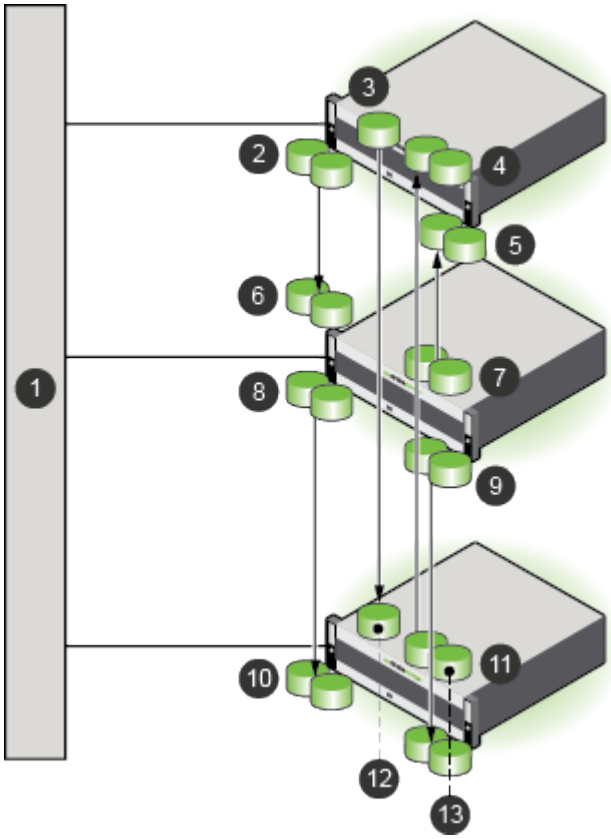

1 Network

2 Multiple volumes that are assigned to the SQL volume collection

3 Single volume that is assigned to the Daily volume collection

4 Replicas of the volumes in the Hourly volume collection

5 Replicas of the volumes in the Datastore1 volume collection

6 Replicas of the volumes in the SQL volume collection

7 Multiple volumes assigned to the Datastore1 volume collection

8 Multiple volumes assigned to the Outlook volume collection

9 Multiple volumes assigned to the Temporary volume collection

10 Replicas of volumes in the Outlook volume collection

11 Multiple volumes that are assigned to the Hourly volume collection

12 Replica of volume in the Daily volume collection

13 Replicas of the volumes in the Temporary volume collection

Replication and Folders

When you create volumes on a downstream partner, you should use a local folder. NimbleOS sends a notification when the folder name matches the upstream name of the volume, so that the upstream volume can be used instead of the pool and folder specified on the replication partner. To preserve the volume organization, Nimble recommends that you manually replicate folder names between partners.

Put Incoming Replicas in a Folder

Use this command to define a partner for an array, and put incoming replicas into a folder on that array. Use the **--match_folder** *yes* option to match the folder for the upstream volume.

Procedure

Create a replication partner, and match the upstream volume folder name with the downstream folder.

partner --create *partner_name* **--folder** *folder_name* **--match_folder** *yes*

Replication Seeding

If you want to replicate a snapshot collection that would take too long over a wide area network (WAN) connection, replication seeding provides an efficient method for creating a snapshot collection to use as a temporary replica. You may need to do this when you first add a replication partner, or if a replication partner needs to be re-initialized due to an accidental deletion or array replacement.

The seeding process replicates a snapshot collection to a temporary array (called the seed) over a high-speed local network (shown as *4* in the image). This third array is transported to the replication site and that same snapshot collection is then replicated from the third array to the designated replication array (shown as *6* in the image).

Once this is completed, the seed array is removed from the system and replications between the upstream replication partner and the downstream replication partner are enabled. From this point forward, the upstream and downstream arrays only need to transfer changed data (shown as *5* in the image).

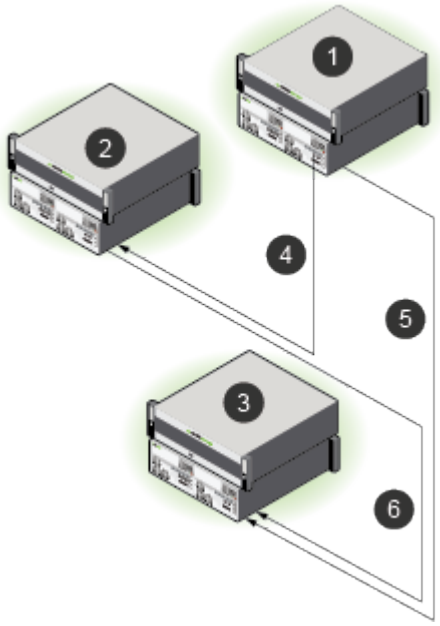

1 Upstream array

2 Seed (temporary) array

3 Downstream array

4 Replication from the upstream to the seed array (one-time)

5 Replication from the upstream to the downstream array (ongoing)

6 Replication from the seed to the downstream array (one-time)

For details about replication seeding, contact support or your sales representative.

Add Replication to a Volume Collection

Before you can set up replication, you need to configure replication partners. Creating replication partners automatically enables two arrays to replicate volume collections that are based on a defined schedule.

To replicate to multiple replication partners, create a separate volume collection for each replication partner. Volume collections can include up to ten schedules. You need to create a schedule and assign it to each replication partner.

Before you begin

You must have a volume collection to create replicas. You cannot create replicas on demand.

Procedure

1. If the volume collection does not include a replication partner in the schedule, edit the volume collection to add a replication partner.

   volcoll --addsched *name* **--schedule** *name* **--replicate_to** *name*

2. Modify each volume to assign the volume to a volume collection.

   vol --assoc *name* **--volcoll***name*

Results

The volumes are now replicated to the designated replication partners according to the assigned volume collections schedule.

Replica Details

Note: To view details about HPE Cloud Volumes (HPE CV), you must access the HPE CV portal.

You can view the details about an on-premises replica on the volumes page for the replica. You can tell that a particular volume is a replica because it has a replica icon in the crumb trail at the top of the page. You can also verify that the **Edit** button is disabled as are the actions in the **More Actions** dropdown menu. The reason for these characteristics is because the replica is owned by the upstream group. When you access the replica from the downstream group, you can only view the replica. You cannot edit it.

The downstream group can claim a replica if the source volume is not replicating, the replication link is broken, and the source volume is no longer part of a volume collection. At this point, the claim option is enabled. After the replica has been claimed, it becomes a regular volume that belongs to the downstream group and all the actions of a regular volume are enabled.

Perform a Volume Collection Handover

You can hand over a volume collection as part of your disaster recovery strategy.

Procedure

1. Perform the handover.

   volcoll --handover *name of volume collection* **--partner** *group_name*

   Note: If you do not want to reverse the direction of replication, add the **--no_reverse** option.

   The new owner is the replication partner group.

2. Verify the handover.

   volcoll --info *name of volume collection*

   Under Schedule, verify that the *Owned by:* and *Replicate to:* have reversed group names.

Example

Example handover.

```
volcoll --handover VolumeCollection5 --partner group26
```

Example verification. Under Schedule, verify that the *Owned by:* and *Replicate to:* column entries have reversed group names

```
volcoll --info VolumeCollection5
```

What to do next

If the replication partners use different access protocols, additional steps are needed to bring the replica online. For example, configuring multi-initiator access is applicable to iSCSI arrays, but is not applicable to Fibre Channel arrays. By default, Fibre Channel allows multi-initiator access, and you cannot disable it. Access control is done through LUN mapping.

Replication Bandwidth Limits

There are two kinds of replication bandwidth limits:

- The overall limit specifies the replication bandwidth for network traffic that originates from an array to all partners
- The per-partner limit specifies the replication bandwidth for network traffic that originates from an array to a specified partner

You can set only one replication bandwidth limit type at a time. If the overall limit is set, the per-partner limits are not allowed for any partner. Conversely, if any partner has per-partner limits set, you cannot set an overall bandwidth limit.

Set Overall Bandwidth Limits for Replication

If you set an overall bandwidth limit for all replications, you cannot set per-partner bandwidth throttling for each replication partner. You can limit bandwidth by using an overall policy or per-partner limits, not both. Bandwidth *limits* are expressed in megabits per second (Mbps) or kilobits per second (Kbps). The default unit is Mbps.

Procedure

Schedule overall bandwidth limits for replication.

group--create_throttle [**--description** *text*]**--days** *days***--at** *time***--until** *time***--bandwidth** *limits*

Example

```
group --create_throttle --days Mon Wed Fri --at 06:00 --until 17:00 --bandwidth
 100Mbps
```

Remove Overall Bandwidth Limits for Replication

If you remove overall bandwidth limits, you can create per-partner policies for each replication partner.

Procedure

Remove the overall bandwidth limits.

group--delete_throttle *id*

> **Note:** Use the **group --info** command to see a list of throttle ID numbers.

Configure Bandwidth Limitations for a Replication Partner

Bandwidth *limits* are expressed in megabits per second (Mbps) or kilobits per second (Kbps). The default unit is Mbps.

> **Note:** If you set per-partner bandwidth limitations, you cannot set an overall bandwidth limit for all replications. You can limit bandwidth by using an overall policy or per-partner limits, not both.

Procedure

Configure bandwidth limitations for a replication partner.

group --create_throttle [**--description** *text*] **--days** *days* **--at** *time* **--until** *time* **--bandwidth** *limits*

Example

Creating a replication bandwidth limitation of 1000 MiB that runs on Monday, Wednesday and Friday, from 6:00 AM until 5:00 PM.

```
Nimble OS $ group --create_throttle --days Mon Wed Fri --at 06:00 --until 17:00
 --bandwidth 1000Mps
```

Modify Per-Partner Replication Bandwidth Limits

When per-partner bandwidth throttling is configured, you can easily modify the schedules. Bandwidth *limits* are expressed in megabits per second (Mbps) or kilobits per second (Kbps). The default unit is Mbps.

Procedure

Modify the per-partner replication bandwidth limits.

partner --edit_throttle *partner_name* [**--description** *text*] [**--id** *number*] [**--days** *days*] [**--at** *time*] [**--until** *time*] [**--bandwidth** *limit*]

Example

Changing the bandwidth limits for partner array34.

```
$ partner --edit_throttle array34 --id 6 --days Mon Wed Fri --at 08:00 --until
 21:00 --bandwidth 500Mbps
```

Security

Security in an HPE Nimble Storage group of arrays includes several features, including role-based access control for user accounts, encryption of data at rest, secure socket layer (SSL) certificates, and secure SMTP.

For information about access to HPE Nimble Storage arrays and volumes and user management, see Access Controls on page 19.

You control access to groups and arrays through the use of user accounts assigned to particular roles. This is known as Role-based Access Control (RBAC). For more information on RBAC, see Role-Based Access Control on page 116.

Role-Based Access Control

Role-based Access Control (RBAC) lets you control access to groups and arrays through the use of user accounts assigned particular roles. The role determines the level of access that a user account is provided. A user account that is assigned to the Administrator role can create and manage accounts for other users.

All users can manage their own passwords and account information. Only users with an Administrator permission level (user role) can add, modify, remove, enable, or disable user accounts.

For more information about permission levels (user roles), see Permission Levels on page 120.

View User Information

All users can view their own account information. Only users with Administrator permission can view the accounts of other users.

Procedure

Type the command for a given user's information.

username --info {**username**}

Example

View user information for Sam:

```
useradmin --info Sam
```

Add a User Account

Each person must have a user account to access and manage an array group. Only users with Administrator permission can add user accounts. The Administrator controls a user's access to the group by assigning a specific role to each user.

At a minimum, each user account must have a:

- username
- full name
- role

Optional. A user account can also have a:

- description
- email address
- inactivity timeout interval

You can include these options when you create the user account or you can add them later.

Procedure

1. At the command prompt, type the command to add a user account.

   useradmin --add {**username**} **--description** {**description text**} **--full_name** {**full name**} **--email_addr** {**valid email address**} **--role** {**administrator** | **poweruser** | **operator** | **guest**} **--inactivity_timeout** {**minutes**}

   - The username must be alphanumeric, 1 to 32 characters, must start with a letter, no spaces. The username is required for user login.
   - (Optional) The description can have from 1 to 255 characters. No hard returns.
   - The full name must be alphanumeric, 1 to 64 characters, must start with a letter, can use dashes, and apostrophes. Underscores are not allowed.

     If the full name has a space, then you must put quotation marks around the full name, for example, "Hunter Smith" Do not use periods.

   - (Optional) An inactivity timeout is specified in minutes. It cannot exceed the inactivity timeout set for the group.

2. At the command prompt, enter a new password.

   {**password**}

   The password must be comprised of alphanumeric characters with a length of 8 to 512 characters. Do not use [ ] & ; ` or spaces. The password is required for user login.

3. Retype the new password.

   {**password**}

Example

Adding a user account for Hunter Smith with operator access:

```
useradmin --add Hunter --description
This_user_has_operator_access_only.
--full_name "Hunter Smith"
--email_addr hsmith@company.com
--role operator
--inactivity_timeout 30
```

Edit a User Account

You can change the full name, role, description, email address, and inactivity timeout interval with this task. You can also add or delete a description, email address, and inactivity timeout interval with this task.

You can use the GUI to set a short interval. You do not need to use the CLI.

Note: In the GUI, the inactivity timeout interval that you set for the group is automatically applied to all user accounts.

Before you begin

You must have Administrator permission to edit a user account.

Procedure

1. Edit a user account.

   useradmin --edit {**username**} **--description** {**description text**} **--full_name** {**full name**} **--email_addr** {**valid email address**} **--role** {**administrator** | **poweruser** | **operator** | **guest**} **--inactivity_timeout** {**minutes**}

2. Verify your changes.

   useradmin --info {**username**}

Example

Editing the user account for Hunter Smith, to change role to administrator.

```
useradmin --edit Hunter --description
This_user_has_operator_access_only.
--full_name "Hunter Smith"
--email_addr hsmith@company.com
--role administrator
--inactivity_timeout 30
```

Change Your Account Information

All users can add, change, or delete their own description and email address. You must log in under your own username to perform this task.

Before you begin

Procedure

Change your account information.

useradmin --edit {**username**} **--description** {**description text**} **--email_addr** {**valid email address**}

The description can have from 1 to 255 characters. Use underscores in place of spaces. No hard returns.

Example

Sam Jones adds a user description and updates his email address:

```
useradmin --edit Sam --description
This_user_has_poweruser_access.
--email_addr samjones@company.com
```

Change Your Account Password

Before you begin

All users can change their own account passwords. You must log in under your own username to change your password.

Procedure

1. Enter the commands for changing your user password.

   useradmin --passwd --user {**your username**}
2. Type your current password.

   {**your current password**}
3. Enter a new password comprised of alphanumeric characters with a length of 8 to 512 characters. Do not use [ ] & ; ` or spaces.

   {**your new password**}
4. Retype the new password.

   {**your new password**}

Example

Sam, as the user, enters the command for a password change.

```
useradmin --passwd --user Sam
```

Reset a User Account Password

When users forget their passwords, perform this task to reset the password.

Before you begin

You must have Administrator permission to reset other users' passwords.

Procedure

1. Enter the command to reset the password for the specified user.
   useradmin --passwd --user {**username**}
2. Type the administrator password.
   {**your Administrator password**}
3. Enter a new password for the user comprised of alphanumeric characters with a length of 8 to 512 characters. Do not use [ ] & ; ` or spaces.
   {**user's new password**}
4. Retype the new password.
   {**user's new password**}
5. Send the user the new password.

Example

Resetting the password for user Hunter Smith:

```
useradmin --passwd --user Hunter
```

Enable a User Account

You can reactivate, or enable, a previously deactivated user account.

Before you begin

You must have Administrator permission to enable user accounts.

Procedure

1. Enable an account for the specified user.
   useradmin --enable {**username**}
2. Verify that the user account was enabled.
   useradmin --list

Example

Enabling an account for user Sam Jones:

```
useradmin --enable Sam
```

Disable a User Account

If needed, you can temporarily suspend a user's access to the array.

Before you begin

You must have Administrator permission to disable user accounts.

Procedure

1. Disable the account for the specified user.
   useradmin --disable {**username**}
2. Verify that the user account was disabled.
   useradmin --list

Example

Disabling the user account for Hunter Smith:

```
useradmin --disable Hunter
```

Remove a User Account

Before you begin

You must have Administrator permission to remove user accounts.

If needed, you can permanently remove a user account. Consider disabling the account rather than removing it.

Procedure

1. Remove the account of the specified user.
   useradmin --remove {**username**}
2. Verify that the account was removed.
   useradmin --list

Example

Removing the user account for Hunter Smith:

```
useradmin --remove Hunter
```

Permission Levels

Each NimbleOS feature has a minimum permission level (user role) that is required to use the feature.

The permission levels (user roles) are summarized here.

Table 10: Summary of Permission Levels (User Roles) and Access

Permission Level (User Role)	Access
Administrator	All actions
Power User	All actions except user management, inactivity timeout, array setup, and array resetup
Operator	Management actions except to delete or remove data
Guest	View information and choose VMware subnets

The following table explains the methods of user access control.

Table 11: How Nimble Interfaces Manage Access

Interface	Method of User Access Control
NimbleOS GUI	Disables or hides unauthorized actions.
NimbleOS CLI	Ignores unauthorized actions and returns a Permission denied message.

The following table lists individual Nimble unit-cs features, the actions associated with each feature, and the minimum permission level (user role) required to perform the action.

Table 12: Features, Actions, and Permission Levels (User Roles)

Feature	Action	Minimum Permission Level (User Role)
Alerts	List / view info	Guest
	Test	Power User
Arrays	List / view info	Guest
	Discover	Power User
	Edit array name	Power User
	Set up / Re-set up	Administrator
	Add	Power User
	Remove	Power User
	Reboot	Power User
	Shut down (halt)	Power User
	Fail over controllers	Power User
Certificates	Regenerate	Administrator
CHAP user	List / view info	Operator
	Create	Operator
	Delete	Operator
	Edit	Operator
Controllers	List / view info	Guest
	Reboot	Power User
	Fail over	Power User
Date and Time	View (local or UTC time)	Guest
	Edit	Power User
Disks	List / view info	Guest
	Add	Power User
	Remove	Power User
DNA	Enable / disable	Power User
	Set secure tunnel	Power User
	Set HTTP proxy	Power User

Feature	Action	Minimum Permission Level (User Role)
Domain name / DNS server for a group	List / view info	Guest
	Add	Power User
Expansion shelves	List / view info	Guest
	Add	Power User
	Activate	Power User
Fibre Channel interfaces	List / view info	Guest
	Edit	Power User
	Update configuration	Power User
Groups	List / view info	Guest
	List limits	Guest
	Edit inactivity timeout	Administrator
	Edit other settings	Power User
	Merge	Administrator
	Create throttle	Power User
	Edit throttle	Power User
	Delete throttle	Power User
	Initiate DNA	Power User
	Validate DNA	Power User
	Unset HTTP proxy	Power User
	Reboot	Power User
	Shut down (halt)	Power User
Help	Show the online help	Guest
Initiator groups	List / view info	Operator
	Create	Operator
	Delete	Operator
	Edit	Operator
	Add initiator	Operator
	Remove initiator	Operator
	Add subnet	Operator
	Remove subnet	Operator
IP addresses	List / view info	Guest
	Add	Power User
	Edit	Power User
	Delete	Power User

Feature	Action	Minimum Permission Level (User Role)
iSNS server for a group	Server IP address	Power User
	Port number	Power User
	Enable / disable	Power User
Migrations	List / view info	Guest
Network configurations	List / view info	Guest
	Create a draft	Power User
	Delete	Power User
	Validate	Power User
	Activate	Power User
	Edit	Power User
Network interface cards (NICs)	List / view info	Guest
	Edit	Power User
	Assign / unassign a subnet	Power User
NTP server for a group	List / view info	Guest
	Add	Power User
Performance policies	List / view info	Guest
	Create	Operator
	Delete	Power User
	Edit	Operator
Pools	List / view info	Guest
	Create	Power User
	Delete	Power User
	Edit	Power User
	Assign / unassign array	Power User
	Merge	Power User
Protection templates	List / view info	Guest
	Create	Operator
	Edit	Operator
	Delete	Operator
	Add a schedule	Operator
	Edit a schedule	Operator
	Delete a schedule	Operator

Feature	Action	Minimum Permission Level (User Role)
Replication partners	List / view info	Guest
	Create	Power User
	Delete	Power User
	Edit	Power User
	Create / edit / delete QoS policy (throttle) settings	Operator
	Pause / resume replication	Operator
	Test connection	Operator
Routes	List / view info	Guest
	Add	Power User
	Edit	Power User
	Delete	Power User
Snapshots	List / view info	Guest
	Delete	Power User
	Edit	Operator
	Set online / offline	Operator
Snapshot collections	List / view info	Guest
	Delete	Power User
	Edit	Operator
SNMP for a group	Edit community string	Power User
	Enable / disable gets	Power User
	Edit responder port number	Power User
	Edit system contact / location	Power User
	Edit trap host IP address	Power User
	Edit trap port number	Power User
	Enable / disable traps	Power User
Software (NimbleOS)	List / view info	Guest
	Download	Power User
	Pre-check	Power User
	Update / upload	Power User
	Resume an update	Power User

Feature	Action	Minimum Permission Level (User Role)
Space reservations, default for a group	Volume reserve	Power User
	Volume quota	Power User
	Volume warning	Power User
	Snapshot reserve	Power User
	Snapshot quota	Power User
	Snapshot warning	Power User
SSH keys	List / view info	Administrator
	Add	Administrator
	Edit	Administrator
	Delete	Administrator
Statistics	Display	Guest
Subnets	List / view info	Guest
	Add	Power User
	Edit	Power User
	Remove	Power User
Time zone	List / view info	Guest
	Change	Power User
User accounts, individual	View own profile	Guest
	Change own password / email address	Guest
User administration	List / view info	Administrator
	Add users	Administrator
	Change passwords	Administrator
	Edit user permissions	Administrator
	Edit timeout interval	Administrator
	Enable / disable users	Administrator
	Remove users	Administrator
User sessions	List / view own info	Guest
	List / view info for other users	Administrator
VMware plugins	List	Guest
	Subnet, choose	Guest
	Register / Unregister	Power User

Feature	Action	Minimum Permission Level (User Role)
Volumes	List / view info	Guest
	Create	Operator
	Delete	Power User
	Edit volume size	Power User
	Edit other volume settings	Operator
	Create / edit / delete ACLs	Operator
	Set online	Operator
	Set offline	Power User
	Take snapshot	Operator
	Restore from snapshot	Power User
	Clone from snapshot	Operator
	Associate with volume collection	Operator
	Disassociate from volume collection	Operator
	Claim partner volume	Power User
	Move to another pool	Power User
Volume collections	List / view info	Guest
	Create	Operator
	Delete	Power User
	Edit	Operator
	Validate	Operator
	Add schedule	Operator
	Edit schedule	Operator
	Delete schedule	Operator
	Take snapshot	Operator
	Restore from snapshot	Power User
	Promote	Power User
	Demote	Power User
	Hand over	Power User
	Stop replication	Power User

Access Control with Active Directory

An Active Directory domain is a collection of objects within a Microsoft Active Directory network. The objects in the domain can include a single user, a group, or a hardware component, such as a computer or printer. Each domain contains a database that stores identity information about the object.

Active Directory relies on mapping Active Directory groups to array roles to determine a user's access. Users are assigned to particular Active Directory groups which are designated with specific array roles. The array roles indicate the level of access permissions that the group members have to perform particular functions.

> **Note:** Active Directory on an HPE array can support up to 100 groups per array with up to 2000 Active Directory users logged in at one time. If you exceed 2000 users, the session for the user who has been logged into the system the longest is terminated.

View Information about the Active Directory Domain

You can use the userauth --info command to view Active Directory domain metadata like the connection status, the organizational unit, and more.

Procedure

List the details about the Active Directory domain.

userauth --info domain_name *domain_name*

The details are displayed in a list as follows:

Active Directory Domain Name : <domain_name>
Organizational Unit : <ou>
Computer name : <computer_name>
Netbios : <netbios>
Status : <status>

Guidelines for Working with Arrays and Active Directory

You can join an array to an Active Directory domain. The following list includes some guidelines to consider when working with arrays and Active Directory.

- When you add an Active Directory user to a group that is authorized to log into the array, you must use the same username and password with the array that you use with all other AD-connected systems in the environment.
- Disabling a user's Active Directory account also disables that user's access to the storage environment.
- Joining an array to a domain creates an Active Directory account for the array on Active Directory. The default Organizational Unit (OU) for the account is Computers. You can create the account under a different OU.
- The default name of the computer account is the first 15 characters of the array group name. You can specify a different computer name if you choose.

  If you use the default name, you must make sure that the first 15 characters of the array group name do not conflict with any other array group name that is in Active Directory. If you duplicate a group name, Active Directory removes the first version of the group name so that the new group name can join Active Directory. Consequently, the group with the duplicated name will not be able to log into the array even though it joined Active Directory first.

  Example:

  1. *group-array-xxxx1* was the first group to join the Active Directory domain *ZZZ*. The default AD machine account name for the array is *group-array-xxx* because the group name truncates after the first 15 characters.
  2. *group-array-xxxx2* was the second group to join the Active Directory domain *ZZZ* but the default AD machine account name would also be *group-array-xxx* because the group name truncates after the first 15 characters.
  3. Upon joining the AD domain, *group-array-xxxx2* replaces the AD machine account from *group-array-xxxx1* with an account for *group-array-xxxx2*. Group users are no longer able to log into *group-array-xxxx1*, although they can now log into *group-array-xxxx2*, which joined later.

- Avoid using special characters in OU and Group names. If you use special characters, they must be preceded by a single backlash, and the entire argument must be inside either single quotation marks (') or double quotation marks ("). For a list of special characters, refer to the Reserved Characters table in Distinguished Names.

- Active Directory administrators can create an account for the array in any OU and then can give storage administrators the privilege to join the domain.
- After an array has joined a domain, you can enable and disable Active Directory authentication without leaving the domain.

Join an Active Directory Domain

Before you begin

You must be a local user with administrative privileges to join an array to an Active Directory domain.

Procedure

1. Join the Active Directory domain.

   userauth --join *domain_name* **--domain_user** *domain* [**--ou** *ou_name*] [**--computer_name** *computer_name*] [**--netbios_name** *netbios_name*]

2. Enter the password at the prompt.
3. (Optional) To verify that the array was added to the Active Directory domain, run the --list command.

Leave an Active Directory Domain

Important: If you leave the domain, all users from that domain lose access to the array and might receive a error message.

Procedure

1. Leave an Active Directory domain to remove access to the Active Directory domain account.

   userauth --leave --domain_user *domain_user_name* [**--force**]

   Note: If you specify force, any error on Active Directory is ignored.

2. Enter the password at the prompt.
3. (Optional) To verify that the array is no longer on the Active Directory domain, run the --list command.

User Authentication and Logon

Active Directory supports the following types of user names for authentication.

- *User_name* (Authenticate with the default domain or as a local user)
- DefaultDomain*User_name* (Authenticate with the default domain)
- TrustedDomain*User_name* (Authenticate with the trusted domain)

Authentication uses the following guidelines:

- If you are authenticating using Active Directory, do not add a group to an array with the group type "Distribution."
- If the array is not a member of an Active Directory domain, then users are authenticated locally on the array. The account must have been created on the array in the **Administration: Users and Groups** dialog.
- If you try to authenticate to an array that is a member of an Active Directory domain, you are authenticated against the Active Directory first.

  Note: Some built-in users, such as root, admin, and nsupport, are always authenticated locally.

  If authentication to Active Directory fails for reasons other than a password failure, the array attempts to authenticate the user locally. If the local account experiences a password failure or the account does not exist locally, authentication fails.

- You can enter a username or a combination of DOMAIN\username. If you do not include DOMAIN, the authentication effort uses the default domain; that is, the domain that the array is a member of.
- The number of repeated failed login attempts allowed depends on the Password lockout setting.

- A successful login provides you with the GUI and CLI roles and capabilities as defined by the group.
- If you lose access to the array, the system response depends on whether you are logged in locally or as an Active Directory user. You might receive an error message, or you might be logged out of the array.
  - When a user is removed from an Active Directory group that has access to the array, the user is no longer able to log into the array. Existing login sessions will continue until the user logs out. This is consistent with the behavior of Windows clients and group memberships.
  - After an Active Directory group is removed from the array, users can no longer log into the group. Existing login sessions will continue until the users log out.

Enable Active Directory Domain Authentication

By default, authentication is enabled after you join the domain. You do not have to enable the Active Directory domain unless it was disabled previously.

Procedure

Enable Active Directory domain authentication on an array.

userauth --enable *domain_name*

The command fails if the status of Active Directory authentication is already in the requested state.

Disable Active Directory Domain Authentication

Important: When Active Directory domain authentication is disabled, users who belong to the Active Directory domain can no longer access the array. For users who are already logged in to the array, any new operation results in an error.

Procedure

Disable the Active Directory domain.

userauth --disable *domain_name*

Active Directory Groups

To use the Active Directory, Active Directory group names must be associated with array user roles. Each Active Directory group can have only one of the following roles assigned to it:

- Administrator
- PowerUser
- Operator
- Guest

If the user belongs to a group that is not associated with any role or if the group is disabled, the user will not be able to login to the array.

If a user belongs to multiple Active Directory groups which have different roles, each time the user executes a CLI command, the group-role mapping with the most restrictive role (the role with the least privileges) is used.

When an administrator makes a change to the group-based RBAC rules, users who are logging in will use the updated roles. For users who are already logged in, they will receive the new privileges for any subsequent operation.

Rules for Array Group Names on the Active Directory

You might want to plan the groups that you want to create on the Active Directory beforehand because group names cannot be changed after the group is created.

It is not recommended that you use special characters in group names.

> **Note:** If you try to add a group (from systems prior to Windows 2000) with a name that contains special characters, the special characters are replaced with underscores (_). Whenever you use that group name, you must use the version of the group name that includes the underscores. The group name that contains special characters is not valid.

You must log in to the array using the pre-Windows 2000 logon name that is shown in the Account tab of the Active Directory server user property page.

Mapping a Group to a Role

To add a group to an Active Directory domain, you must map an Active Directory group to a specific role.

If the user belongs to a group that is not associated with any role or if the group is disabled, the user will not be able to login to the array

> **Note:** For users of systems prior to Windows 2000: If you try to add a group with a name that contains special characters, the special characters are replaced with underscores (_). Whenever you use that group name, you must use the version of the group name that includes the underscores. The group name that contains special characters is not valid.

Procedure

1. Map a group to a particular role.

   userauth --add_group *group_name* **--domain** *domain_name* [**--role** *role*] [**--description** *description*] [**--inactivity_timeout** *minutes*]

   If the role is not specified, the guest role is assigned to the array group.

   If the inactivity_timeout is omitted, the inactivity timeout for the array group is used.

2. Run the following command to verify that the process completed successfully.

   userauth --list_group --domain *domain_name*

Enable an Active Directory Group

Procedure

Enable the Active Directory group.

userauth --enable_group *group_name* **--domain** *domain_name*

Remove an Active Directory Group

> **Note:** After a group is removed from the array, users can continue to log into the group and perform actions for approximately five minutes. After five minutes, the user can no longer access the group.

Procedure

1. Remove a group from the Active Directory domain.

   userauth --remove_group *group_name* **--domain** *domain_name*

2. (Optional) Verify that the group was dissociated from the role.

   userauth --info *domain_name*

Disable an Active Directory Group

> **Note:** If a group is disabled then users will not be able to log in to Active Directory.

Procedure

Disable the Active Directory group.

userauth --disable_group *group_name* **--domain** *domain_name*

Edit Active Directory Group Information

You can edit the following information in an Active Directory group:

- Group role
- Group description
- Inactivity timeout

> **Note:** If the group name is changed, you must remove the group mapping on the array and add a new group mapping with the new group name.

Procedure

Edit the Active Directory group.

> **Note:** You cannot change an Active Directory group name after the group is created.

userauth --edit_group *group_name* **--domain** *domain_name* [**--role** *role*] [**--description** *description*] [**--inactivity_timeout** *minutes*]

Troubleshooting the Active Directory

When you have issues with your array on the Active Directory, you can run some of the following commands to try to fix the issue.

Command String	Function
userauth --test_group *group_name* **-- domain** *domain_name*	Tests whether the group exists in the Active Directory domain
userauth --test_user *user_name* **-- domain** *domain_name*	Tests whether the user belongs to the Active Directory domain
userauth --info	Provides view of the current Active Directory status
userauth --disable *domain_name* **userauth --enable** *domain_name*	If you make any changes on the Active Directory server that are related to the Active Directory configuration on the array, you might need to disable Active Directory authentication and re-enable it to clean up the cache.

For complete syntax on these commands, refer to the *Nimble Storage Command Reference*.

CHAP Authentication

As the name implies, Challenge-Handshake Authentication Protocol (CHAP) uses a challenge-response mechanism to authenticate iSCSI initiators. A shared "secret," or password, let the system verify that the iSCSI initiator is who it claims to be and is authorized to access the volume.

Before you can use CHAP authentication, set up the CHAP secret on the volume and on the iSCSI initiator. CHAP secrets must be between 12 and 16 characters long. For the best security, the secret should be random letters and numbers, not a word

that could be guessed. If your iSCSI initiator imposes further restrictions on the CHAP secret, you must adhere to these stricter regulations.

When creating a CHAP secret, adhere to the strictest regulations: 12-16 characters containing no spaces or the special characters (' " ` ). The CHAP user name should not contain characters such as : ~ ! @ # $ ^ & ( ) + [ ] {} * ; : ' " ., % | < > ? / \ = `.

Create a CHAP Account

CHAP (Challenge Handshake Authentication Protocol) users share a "secret." This CHAP secret is a word, phrase, or series of characters that both the array and the initiator know. The array will only allow access to those iSCSI initiators who respond with the correct secret.

Procedure

Create a CHAP user

chapuser [**--create** *name*] [**--password** *shared_password*]

Assign a CHAP User to a Volume

The CHAP user must be created before it can be assigned to a volume. Multiple volumes can be assigned to the same CHAP user.

Procedure

Assign a CHAP user. Specify whether to add the ACL to a volume, a snapshot, or both.

vol --addacl *volume_name*

Modify a CHAP User

For increased security, you may want to change the CHAP secret at regular intervals, or if you suspect that an unauthorized computer has accidentally gained access to the array.

> **Note:** If you change a CHAP secret, all volumes protected with this CHAP account will be inaccessible until the corresponding iSCSI connections are changed and synchronize with the new CHAP secret. Consider then when you determine what time to make these changes.

Procedure

Edit CHAP user information. Enter at least one parameter to be changed (name, description, or password).

chapuser [**--edit** *name*]

Delete a CHAP User

Delete CHAP users that are no longer needed.

Procedure

Delete a CHAP user.

chapuser --delete *name*

Login Banner

By default, a login banner displays in the CLI interface for all controllers in an array group. However, the banner can be configured to not be displayed (deleted). It can also be configured to be displayed either before prompting for user's credentials, or after user authentication. By default, the login banner is displayed after user authentication.

Important: You must have Administrator privileges to configure the login banner.

The login banner has a factory default login banner message, but the message can be edited to suit your specific requirements. The message is restricted to 2,048 ASCII printable characters with support for newline. International characters are not supported. (The official Department of Defense [DoD] banner character count is about 1,200 characters.) An edited banner message can be reset to the factory default message.

Edit the Login Banner

Before you begin

Ensure that you have Administrator privileges.

Procedure

1. Enter the following command to edit the login banner.

   group --edit [**--login_banner**] [**--login_banner_after_auth** {**yes** | **no**}] [**--force**]

2. After you type the command and press Enter, type the banner message, then type ctrl + D to save the message.

   Note: The login banner will be deleted if you create an empty banner message (no characters typed). To delete the login banner use the --force option when entering the command.

Example

Editing the login banner message and setting the login banner to be displayed before user authentication.

```
$ group --edit --login_banner --login_banner_after_auth no
Please enter the banner message below followed by ^D:
*********************************************
*                                           *
* SECURITY WARNING: AUTHORIZED PERSONNEL ONLY *
*                                           *
*********************************************
$
```

Deleting a login banner

```
$ group --edit --login_banner -force
Please enter the banner message below followed by ^D:
$
```

Show the Login Banner

Procedure

Enter the following command to show the login banner.

group --show_login_banner

Example

Showing the login banner.

```
$ group --show_login_banner
USAGE WARNING

This is a private system. This system is provided only for authorized use.
Unauthorized or improper use of this system may result in civil claims
and/or criminal charges. The array owner may monitor the system for all
```

```
lawful purposes, including but not limited to ensuring that access is
authorized and for other security reasons. Use of this system constitutes
consent to the array owner for monitoring of this system.

Administrators should ensure that this system is protected by a firewall
and implements all security procedures itemized in the documentation.
$
```

Reset the Login Banner

Procedure

1. Enter the following command to reset the login banner message to the factory default message.

   group --reset_login_banner

2. (Optional) Verify that the login banner message has been reset to the factory default message

   group --show_login_banner

Example

Resetting a custom login banner message to the factory default message

```
$ group --show_login_banner
*********************************************
*                                           *
* SECURITY WARNING: AUTHORIZED PERSONNEL ONLY *
*                                           *
*********************************************
$ group --reset_login_banner
$ group --show_login_banner
USAGE WARNING

This is a private system. This system is provided only for authorized use.
Unauthorized or improper use of this system may result in civil claims
and/or criminal charges. The array owner may monitor the system for all
lawful purposes, including but not limited to ensuring that access is
authorized and for other security reasons. Use of this system constitutes
consent to the array owner for monitoring of this system.

Administrators should ensure that this system is protected by a firewall
and implements all security procedures itemized in the documentation.
$
```

Encryption of Data at Rest

You can enable encryption at the group level or at the volume level as required for each group of arrays in your environment. Before you can create encrypted volumes, you must perform an initialization step that creates the master key. The master key protects the keys that are used to encrypt volume data. The master key is protected by a passphrase that is specified when creating the master key. At times, it will be necessary to enter the passphrase to enable access to encrypted volumes.

The encryption state of a volume is established when the volume is created, and cannot be changed afterward. Cloned volumes inherit the encryption state of their parent. The group configuration contains a default encryption default setting, where you can either enable or disable AES-256-XTS encryption. (The AES-256-XTS encryption algorithm is specifically designed for use in encrypting block storage.) The group configuration also contains an encryption scope setting, which specifies where and how to apply the encryption default setting. You can force the encryption default setting to be applied to all new volumes in the group, or allow overriding the encryption default setting on a per-volume basis.

The group configuration contains an encryption mode setting that defines behavior on system restarts. The value can be set to "secure" or "available." In secure mode, the encryption passphrase must be entered every time the group leader array is

restarted to unlock the master key. In most cases, available mode stores enough information in non-volatile memory to recover the master key without entering the passphrase. The information is not stored on disk. Available mode is provided for convenience in situations where the physical security of the array is unlikely to be compromised.

Important: Even though available mode significantly reduces the number of times you must enter a passphrase when a group leader array restarts, it does not guarantee that you will never have to enter a passphrase after a restart. There are certain scenarios where you would still have to specify a passphrase while in available mode to access encrypted data, including:

- Controller upgrade: If array controllers are being upgraded to a newer model, you must enter a passphrase. While data is recovered from the non-volatile memory, access to encrypted volumes is not.
- NVRAM loss: In the rare case where non-volatile memory (NVRAM) is lost, you must enter a passphrase to access encrypted volumes. Older arrays (CS2xx and CS4xx series) that remain powered off for a long time could lose NVRAM as a result of battery discharge.

CAUTION:

- If you lose the passphrase for the master key or access to the external key manager, data in encrypted volumes cannot be retrieved. Store the passphrase in a secure, accessible place.
- If your encryption requirement changes after creating a volume, you cannot change its encryption status. You can create a new volume with the encryption status that you need, and migrate the data to the new volume.
- Performance might be slow when accessing encrypted volumes from the CS210 or CS215; however the performance impact due to encryption will be less severe on the CS235 arrays.

Enable Encryption

Beginning with version 6.0, you have the option of using a passphrase for local key management or using external key management for your encryption keys. The use of encryption involves using keys to encrypt volume data. Two important points to remember are:

- If you lose the passphrase for the master encryption key or if you lose access to the external key manager, data in the encrypted volumes cannot be retrieved.
- Once it has been set, the encryption status of a volume cannot be changed.

Before you begin

You must have Administrator privileges to change the encryption configuration.

To ensure that you are aware of the requirements for encrypting volumes, read the information in Encryption of Data at Rest on page 134.

Procedure

1. Create the master key or set up the external key manager.

   encryptkey --create_master

2. Enter a new passphrase composed of any printable characters with a length of between 8 and 64 characters.
3. Retype the new passphrase.
4. (Optional) Specify the group encryption settings.

   group --edit --encryption_cipher {*aes-256-xts* | *none*} **--encryption_scope** {*group* | *volume*} **--encryption_mode** {*available* | *secure*}

   The group encryption settings are applied to the volumes that you create from this point forward. The settings are not applied to existing volumes.

5. Create a volume using the encryption settings that you need and that are valid based on the group encryption settings.

   vol --create *volume_name* **--size** *mebibytes* **--encryption_cipher** {*aes-256-xts* | *none*}

Note: After volume creation, encryption on that volume cannot be changed.

Example

Enabling encryption using the default group encryption settings (encryption_cipher = aes-256-xts, encryption_scope = group, and encryption_mode = available). Here you have the option of using the encryptkey --create_master command to enable encryption with local key management or setting up an external key manager as described in the next section.

```
$ encryptkey --create_master
Enter new passphrase:
Retype new passphrase:
```

Creating a volume with encryption when encryption is enabled for the group.

```
$ vol --create finance --size 1000000 --encryption_cipher aes-256-xts
```

Creating an unencrypted volume when encryption is enabled for the group.

```
$ group --edit --encryption_scope volume
$ vol --create facilities --size 500000 --encryption_cipher none
```

Secure Sockets Layer Certificates

To establish a secure connection with a website or other server, the server presents a certificate to authenticate its identity. Certificates are an important component of Secure Sockets Layer (SSL) because they prevent others from impersonating a secure website or other server.

There are three types of certificates.

- **Array certificate chain:** Generated when the array is first started.
- **Group certificate chain:** Generated when the array is configured as a group leader.
- **Custom certificate chain:** (SSL Certificate) Either a self-signed certificate or a certificate generated by exporting the CSR then signing and importing the root CA and signed certificates. This is the most secure type of certificate.

An SSL certificate is an electronic document that verifies ownership of a public key and ensures the identity of your server, providing greater security of online interactions. The certificate includes the following information:

- Information about the key
- The identity of its owner
- The digital signature of a trusted certificate authority (CA)

The digital signature verifies that a trusted third party (the CA) has authenticated the identity of the organization that owns the key and has verified that the contents of the certificate are correct.

If the signature is valid, and the person examining the certificate trusts the signer, then they know that it is safe to use that key to communicate with its owner.

To get an SSL certificate, you must create a Certificate Signing Request (CSR). Then, you send the CSR data file to the CA and the response that you receive from the CA is your SSL certificate. This SSL certificate is the intermediate chain public key and you import the key through the CLI.

After you receive the certificate and install it on your server, the identity of your server can be authenticated.

Create and Import a Custom-Signed Certificate

You use the cert --gen command to create either a self-signed certificate or a certificate signing request (CSR).

If you use the custom argument, the command generates a self-signed certificate of the type indicated with the options that you specify. If you use the custom-csr argument, the command generates a CSR with the options that you specify. After the certificate is signed, you can import it as a custom certificate.

For more information about the cert command, refer to the Nimble Storage Command Reference.

Procedure

1. Create the certificate or CSR:

   cert --**gen** {**array** | **group** | **custom** | **custom-csr**} [**--subject** text] [**--dnslist** text] [**--iplist** text] [**--num_days** text] [**--check**] [**--force**]

   NimbleOS $ cert --gen custom-csr --subject '/C=US/ST=CA/L=San Jose/O=Nimble Storage/OU=Engineering/CN=AF106656' --dnslist group-kp-vma.nimblestorage.com,kp-vma.nimblestorage.com

   Then

2. Cut and paste the certificate request output into a CA signing request

   NimbleOS $ cert --import custom-ca

3. Cut and paste the CA certificate and the output from the signing into the command inputs below:

   NimbleOS $ cert --import custom

Delete a Certificate

You can only delete a custom certificate because the array and group certificates are automatically generated.

Before you begin
You must have a custom certificate installed on your server.

Procedure

1. Delete a custom certificate.

   cert --delete *custom*

2. (Optional) Enter the following command to verify that the certificate was deleted.

   cert --list

Create a Custom Certificate Chain

To verify that the array can trust the certificate, you must present the entire chain of certificate authorities the array. The certificate chain must include the root certificate and every subordinate certificate (there might be more than one subordinate certificate).

Before you begin
It is important that you paste each subordinate certificate into the certificate chain in order.

Procedure

1. Gather your certificates from each CA server in the chain.
2. Open the root CA certificate as a text file.
3. Copy the contents of the root CA certificate into notepad.

   Make sure to include the begin certificate tags and the end certificate tags for each certificate.

   -----BEGIN CERTIFICATE-----

   -----END CERTIFICATE-----

4. After the last line of the root CA certificate, enter a carrige return.
5. Paste the contents of the first subordinate CA certificate.
6. After the last line of the first subordinate CA certificate, enter a carrige return.

7. Repeat steps 5 and 6 for each additional subordinate CA certificate.
   This example shows one root CA certificate and one subordinate CA certificate. Your certificate chain might include more than one subordinate CA.

   -----BEGIN CERTIFICATE-----

   MIIGXjCCBEagAwIBAgIQZe4IxsgABbZCTj2L0apZATANBgkqhkiG9w0BAQsFADBS

   MRUwEwYKCZImiZPyLGQBGRYFbG9jYWwxFjAUBgoJkiaJk/IsZAEZFgZjYXJzb24x

   ITAfBgNVBAMTGGNhcnNvbi1UQ0FSU09OLVJPT1RDQS1DQTAeFw0xNzEwMjYxMjE1

   MTFaFw0xOTEwMjYxMjI1MDlaMFIxFTATBgoJkiaJk/IsZAEZFgVsb2NhbDEWMBQG

   CgmSJomT8ixkARkWBmNhcnNvbjEhMB8GA1UEAxMYY2Fyc29uLVRDQVJTT04tUk9P

   VENBLUNBMIICIjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEAzL+9d87ypqEo

   rqSmcqXLRy1/rcLFVE5ozxPI1vfwUqLaLQC6SKPow6xD0ZQrvtryehrZQ91sjsSr

   9bx/7f9NloMRTeRzEoH2wv0rfUijGV78B2lNMYPqIxmTOeefKE9DIlWDyT3R1kaj

   RFrbuJFVCJVCfS2Jwrw8FQ56AzL9CbEFXYi4ITCNdd4ZsICll6Qwxg+8scVFFCta

   527TXQB6wfCtomk4FW94/q//TgPLbuO/RwwMTkEaAa2Nipl+qLAh3rpTvrprGAtj

   A7fAeEdrcK3hPlDz9rRqTL24xd6qh5MaUNCZhxxLZg3auu6FPcSz5ykoGD81QwaE

   nC8wFbDJE6/5HdbAtboxIg+HelhT6rsp0b0AzVPglM47pHNC37MUlFD1qBO+pA+d

   5Ox0Lj9/I1xrVSHBHESu0cYWqtQN728WPg4ZxivL7jzo3RyCwp+9Nw6JzxwBzLdz

   MiDShVYB+A6nwppOCopxXhistrOweGhflucmz++bq1KrH7DsNa5gmGGWTMytWaYB

   QKe5EDVQaK9i63fzBhNWaZiSYJhRZRGNOYmY/VOwhFlSR6jjcwOnnWhWM3NnkdSG

   fT18mpclqbhFX8FIlyXyazqCmVCd8e3urpHzI10x6JJMQStqj4APF7qsEXWW46vJ

   hjAPBgNVHRMBAf8EBTADAQH/MB0GA1UdDgQWBBT0ePFtg+E4caBFCkvHsbFklaRp

   xzASBgkrBgEEAYI3FQEEBQIDAQABMIGxBgNVHR8EgakwgaYwgaOggaCggZ2GT2Zp

   bGU6Ly8vL3RjYXJzb24tcm9vdGNhLmNhcnNvbi5sb2NhbC9DZXJ0RW5yb2xsL2Nh

   cnNvbi1UQ0FSU09OLVJPT1RDQS1DQSgxKS5jcmyGSmh0dHA6Ly90Y2Fyc29uLXN1

   YmNhLmNhcnNvbi5sb2NhbC9jZXJ0ZGF0YS9jYXJzb24tVENBUlNPTi1ST09UQ0Et

   Q0EoMSkuY3JsMCMGCSsGAQQBgjcVAgQWBBTtwFmmGeyZ0Ap+0bJUol08Ul6xyz

   BgkqhkiG9w0BAQsFAAOCAgEAdcgfjxG+dj0Frw0x9vCPaCHwsP31qA/br7VRQQXE

   p9xQ96xvO7gyn/UNqApxl3Jlb11m5S0pJC1QcHt6s6HvVrDaNJTQCpTn57qrIM/t

   zOCExe1u5qvEE0l7ruk93e10EjTrDh1NZ5YSyMyButxWzSYevQz9c/j4SKxPms7P

   GTTNvCXQrzVPZ2Olhwz7jKw63CnmLjMZvwcXU1++mp0W1lCdlUwDW6JoZqdHk6Pn

   h19tI2jHGEuaoXy28dhFaA9vslCYswE0oT8w1svJSYTWPqUZhnZc60N75LWfPVHA

   pM1COjGiVo3bezYpWUhMliCsnLx33w5mkmoIkReU27+x1aE5CG0cpeXWd1KoGU7k

   YkqyTpFVy2hnDipbAuumgGmlR88kV0yOZZkGT9Oje9oV/hanyO68VMVkklegY5kp

   A6n66hJIGg3rUWIVIQArHOsaPnml5I3n06Yt04SPNG8S6Pq0BZDbQ8nNjLRCpNMq

   gRHgStXsZzcxna3QQvcdCzLe5H2xo3bf90lZD6firaOQkO1MSDZn6CvGhcFn8fZH

   n8xWKRJjC8ppOaSvKCcvSBulXcuQEf6CE3QMhDxXUUbWxj2mkzMljYFv7dIhWD9R

   x7/4Y42HzDaGh+g6QQzSlx9tYsltx0Fl7BKBWI4D7PqcsO3eTKze+TWoks1Cj6b1

   BtOotfVPbJR59MPQV7typNgpLAWHoY8CAwEAAaOCAS4wggEqMAsGA1UdDwQEAwIB

1hw=

-----END CERTIFICATE-----

-----BEGIN CERTIFICATE-----

MIIHmzCCBYOgAwIBAgITYgAAAAiI2D6wVb9OUAABAAAACDANBgkqhkiG9w0BAQsF

ADBSMRUwEwYKCZImiZPyLGQBGRYFbG9jYWwxFjAUBgoJkiaJk/IsZAEZFgZjYXJz

b24xITAfBgNVBAMTGGNhcnNvbi1UQ0FSU09OLVJPT1RDQS1DQTAeFw0xNzEwMjYx

MjE4MjBaFw0xODEwMjYxMjI4MjBaMFExFTATBgoJkiaJk/IsZAEZFgVsb2NhbDEW

MBQGCgmSJomT8ixkARkWBmNhcnNvbjEgMB4GA1UEAxMXY2Fyc29uLVRDQVJTT04t

U1VCQ0EtQ0EwggIiMA0GCSqGSIb3DQEBAQUAA4ICDwAwggIKAoICAQCs/ZyVt26L

I4aalhaCJeIaw9ESUN8mfoqC7S5OY1eLFH3u8y2cLXgs1DRi8Qbh9/E23++zHk+1

v+L0eGTeQVDNqEEHl5o4mSwrZxfZw/HNOrHWVZcbLcxrbhQQeQxyAO8Jxys4tiaf

cKvRAiJMmpiByn16TAmDxj/tUKd3Ggr70i5rYv1VSgM6e4wtXtPvYYYwVyriFt4e

GbJAjKDjZLkgf/bKhiZjTCEvQzdz1RjFGe975mxGtrdY9sE1P/iaQekKVFQp76p0

3oUvpQRW8INA++2dYb69RdazkD82etPzdY+XXLvF3PdSKIvNz2WXfrfoJRIjlXJ2

qR30UWjwvRn8vAnaZeU8cS9Bv0ikWGNzONuGPe/4q4TmN4NMX2/wIRtBN3XHDx

6EIQDt+ks3QJdq/mNmF0WFNwq6gkN53JR/H1WPDXWoRTIbFH/6kzST5MXTQA4gUF

Uu9vycpn7w91pgiWLZHN0m8FqHF0Uf4hSHoGxr0AUs/oV5yXGa/sC0YbK8EZmvlg

SXgt1sYRNqV8TFt4F/JjPjqwobQTvClBvd2l4oM1l0pjqSU28vASbvyUbA1Gl7GR

jEvYXB8B531zHJKIu0dy6JstNpbwDqzdOTUvJQvVqThRU5NKXsa3O/jRM0PMK/Sx

jzSfhBrX+5PQBm3WhuDw4LkVwk7HUigdvwIDAQABo4ICaTCCAmUwEgYJKwYBBA

NxUBBAUCAwEAATAjBgkrBgEEAYI3FQIEFgQULcg+TBDHgDtkggANblRuzdYt49kw

HQYDVR0OBBYEFEAnMn2zL8JY0d4KFxQaR2HbRjSqMBkGCSsGAQQBgjcUAgQMHgo

MBaAFPR48W2D4ThxoEUKS8exsWSVpGnHMIGxBgNVHR8EgakwgaYwgaOggaCggZ2G

T2ZpbGU6Ly8vL3RjYXJzb24tcm9vdGNhLmNhcnNvbi5sb2NhbC9DZXJ0RW5yb2xs

L2NhcnNvbi1UQ0FSU09OLVJPT1RDQS1DQSgxKS5jcmyGSmh0dHA6Ly90Y2Fyc29u

LXN1YmNhLmNhcnNvbi5sb2NhbC9jZXJ0ZGF0YS9jYXJzb24tVENBUlNPTi1ST09U

Q0EtQ0EoMSkuY3JsMIH8BggrBgEFBQcBAQSB7zCB7DB3BggrBgEFBQcwAoZrZmls

ZTovLy8vdGNhcnNvbi1yb290Y2EuY2Fyc29uLmxvY2FsL0NlcnRFbnJvbGwvdGNh

cnNvbi1yb290Y2EuY2Fyc29uLmxvY2FsX2NhcnNvbi1UQ0FSU09OLVJPT1RDQS1D

QSgxKS5jcnQwcQYIKwYBBQUHMAKGZWh0dHA6Ly90Y2Fyc29uLXN1YmNhLmNhc

bi5sb2NhbC9jZXJ0ZGF0YS90Y2Fyc29uLXJvb3RjYS5jYXJzb24ubG9jYWxjYXJz

b24tVENBUlNPTi1ST09UQ0EtQ0EoMSkuY3J0MA0GCSqGSIb3DQEBCwUAA4ICAQCl

N3TRmkDghlwYg/Xhd9KJjcOBUG6ZuhxqywKk1hxOA1a57+3eri4GrXdy7/kapsIF

nMKrlaHR9yAC8W/YZuCrWDmoiND1OyB+jyG5sCdnB6eEFkky3t1Q1H9E1W2hcHPO

uOc/faWtuxMAtzJilLUQVpUg2GP0yoYoF4xNp3X7MmXb3xhkSS+ZXW69VEFKBFxW

E+ovAWNFZndy9V/nlDQiXrobN5FUoYAFTujkfxXM5folc46sm2MB55tC16idsHs4

nEk9N/zyCBEbFx+n1GuYTo6RlKjaPqgsPBtpWAyc0HXXSKZ+KhkG0rmLga/UbvTM

vetQ3Dvp8WgvtYg37PJUPZo9yXH5Cfr3AVLBZLRMAwUNWdY5y68oi7lh474aO2CX

ODMLcktoAYqHIXNxZttYeez6OUaU/jJ8oZLDtVCfCKCHATVwyqfCi/tE6yt3SYrv

Br3Jucp0PZlSdhHchu8YtGeL/0U61DeH/l2HoTpeq71q1rHc4nW13GV87riHCf6g

U26n7NbzSBCRPzhs3CtM/qDW8ezkbN0la2CUKUbVw4oqVgay1oapdUkrNHC7W/Xq

T+PNkxKWAV7uDaPLk2HGlaur38QJetL2HU69fZ4pU4e06MoX6g0ROSMzU18eUGkq

UwB1AGIAQwBBMAsGA1UdDwQEAwIBhjAPBgNVHRMBAf8EBTADAQH/MB8GA1Ud

p8ZWLDnaa0ZPQBtkp/n1i49yd3qotx42mKmV9JA7tA==

-----END CERTIFICATE-----

8. On the array, import the custom CA.

   cert --import *custom-ca*

   The array returns the following message.

   Please enter certificate in PEM format followed by ^D:

9. Paste the entire certificate chain (as in the example) into the array.
10. Press enter to add a carriage return after the last end certificate tag.

    Note: If you do not add a carriage return, the process will fail.

    If the process is successful, the array OS command prompt returns.

Specify a Certificate Chain to Use to Authenticate HTTPS and API Services

Before you begin

You have a created custom certificate chain. For more information on creating a certificate chain, see Create a Custom Certificate Chain on page 137.

Procedure

1. Specify which certificate chain to use to access the web.

   The argument indicates which of the three certificates to specify.

   cert --use {*array* | *group* | *custom*} [**--https**] [**--apis**]

2. (Optional) Enter the following command to verify that the correct certificate chain was specified.

   cert --list

Multihost Access (MPIO)

The Nimble array supports multihost access also known as multipath I/O (MPIO). When an initiator connects to a target, the access control records do not automatically prevent multiple initiators to connect. As long as the access control record limitations are met, the initiator can connect.

In some environments, you may need multiple initiators to access a target. These conditions include the following:

- A virtual server that manages multiple connections
- An environment in which initiators on the same computer do not use the same IQN
- An environment that uses a Distributed Lock Manager

Using MPIO

Note: Ensure that you have an active iSCSI connection before installing MPIO. Not having an active connection before installing MPIO causes the Add support for iSCSI devices feature to be unavailable.

Install an MPIO product on the system that is accessing the array. MPIO requires multiple network adapters dedicated to the iSCSI task. When connecting your iSCSI initiators, select Properties and click MPIO.

MPIO determines which paths to a device are in an active state and can be used for load balancing. The load balancing policy (Least Queue Depth is recommended by HPE) is set in the DSM. This policy determines how the I/O requests are actually routed.

MPIO for Windows

For information on installing and configuring MPIO on Windows, refer to the *Windows Integration Guide*. It was based on installing MPIO onto a Windows 2008 Server. MPIO is an optional component with Windows 2008 Server. The process is similar on a Windows 2003 server after obtaining the MPIO component.

MPIO for Linux

For information on installing and configuring MPIO on a Linux-based system, refer to the *Deployment Consideration for Linux on Fibre Channel* and *Deployment Considerations for Linux on iSCSI*.

Secure SMTP

You can configure simple or secure Simple Mail Transfer Protocol (SMTP) to send alerts from groups to external servers. Alerts are identifiers about specific actions that occur on a group of arrays.

Prior to version 2.3, you could configure only a simple (or regular) SMTP relay of email alerts.

In version 2.3 and later, you can configure either a regular or a secure SMTP relay.

Configure Email Alerts

You can configure email alerts to use regular or secure Simple Mail Transfer Protocol (SMTP) processing, depending on which mode is appropriate for your environment.

You might choose to use regular SMTP for email alerts if you have an SMTP server installed on your network that accepts email messages from external parties. You might choose to use secure SMTP if you have an SMTP server installed on your network, but prefer to disallow anonymous relays, or if you do not have an internal email server because you implemented cloud-based email, such as Office 365.

You can configure email alerts differently for each group of arrays. You must have at least Power User permission to configure SMTP-based email alerts.

Using the following command options you can modify group settings, change the hostname or IP address of the SMTP server or its port number, allow use of SMTP authentication and encryption (with options to change the authentication password and enable the level of encryption, when enabled), change the email address for sending and user receiving email alerts, enable email to be sent to support, and set the minimum event data alert level at which alerts are sent.

For a full list of all group --edit options, refer to the *Command Reference*.

Procedure

Configure email alerts.

group--edit [**--smtp_server** *smtp server*] [**--smtp_port** *smtp port*] [**--smtp_auth** {**yes** | **no**}] [**--smtp_username** *username*] [**--smtp_encrypt_type** {**none** | **starttls** | **ssl**}] [**--smtp_from_addr** *email addr*] [**--smtp_to_addr** *email addr*] [**--send_event_data** {**yes** | **no**}] [**--alert_level** {**info** | **warning** | **critical**}]

> **Note:** When you edit SMTP, you must always include the --smtp_username, --smtp_encrypt_type, and password to prevent the operation from failing. If you are using authorization, you must also include --smtp_auth yes. You do not enter the password on the command line. After you run the command, you will be asked to enter the password. Providing the password this way avoids having to display it in plain text.

Note: The syntax for configuring SMTP email to send alert messages for critical alerts to multiple email addresses is, for example:

```
--smtp_to_addr Fred@here.com --smtp_to_addr George@here.com
```

Example

Configuring SMTP email to send alert messages for critical alerts from user Joe Black to support.

```
$ group --edit --smtp_auth yes --smtp_username JoeBlack@mycompany.com --
smtp_encrypt_type ssl --smtp_from_addr Joe.Black@customer.com --smtp_to_addr
support@nimblestorage.com --alert_level critical
```

Monitoring Your Arrays

A Nimble array needs little ongoing maintenance after it is installed and configured. Even though the system runs without extensive administrative activity, it is a good idea to monitor the system regularly to make sure that everything is working correctly.

You can choose options from the NimbleOS **Monitor** menu to monitor the array in real time. Monitoring lets you track system trends and proactively ensure that no bottlenecks occur. Nimble's intuitive monitoring system lets you see space usage and performance at a glance.

Several monitoring options use a set of common controls. When you monitor capacity, performance, interfaces, connections, the audit log, or replication, you can specify the time interval of interest to you. Choose from:

- **Real-time**, which is useful for monitoring real-time activity
- Last 3 minutes (**3M**), which is useful for monitoring very recent activity
- Last 60 minutes (**1H**), which is useful to determine whether an activity is a temporary or recurring condition
- Last 24 hours (**1D**), which is useful for tracking activity patterns for the day
- Last 7 days (**1W**), which is useful for tracking activity patterns for the week
- Last 30 days (**1M**), which is useful for tracking activity patterns for the month
- Custom, which lets you specify a time interval of interest

The selected interval determines how much data is shown. The longer the time interval, the more compressed the data appears in the graph. Use longer time intervals for tracking trends that you can use for purchasing estimates and capital expense projections.

When you monitor **Capacity** or **Performance**, you can also select one or all volumes to include in the data collection. By default, all volumes are included. However, you can limit the display to a specific volume.

The other monitoring options do not have the **Real-Time** and **Volume** common controls. Those pages provide other ways to track activity patterns.

Monitor Space Usage

You can monitor space usage to track space-heavy applications and usage trends.

The **Space** page displays volume and snapshot usage, unused reserved space, and free space for one or all volumes on an array. The upper part of the page provides information about the space usage for the array. The lower part provides information about space usage that is based on time and volume selections.

Procedure

Monitor space usage using the CLI.

group --info *group name*

> **Note:** In the CLI, the information returned is current at the moment you run the command. The command output does not change as the system automatically refreshes the space usage.

Monitor Performance

You can monitor the performance levels of all volumes on an array or a specific volume.

Procedure

Monitor the performance of your volumes.

stats --perf *volume_name* [**--latency**] [**--iosize**]

Monitor Interface Traffic

You can monitor throughput for interfaces to determine whether traffic is balanced appropriately.

Procedure

1. Monitor interface traffic.

   stats --array *array_name* [**--from** *time*] **--to** [*time*] [**--duration** *time_interval*]

   stats --array array5 --from 08:00 --to 15:00
2. (Option) To view more targeted interface stats, run the command associated with your array type.

Option	Description
On an iSCSI array	**stats --net** {*all* | *specific_nic_name*} Use **nic** --list to get values for *specific nic name*.
On a Fibre Channel array	**stats --fc** {*all* | *specific interface name* } Use **fc**--list to get values for *specific interface name*.

Monitor Replication

You can monitor the lag time of replications when sending data to a partner or receiving data from a partner.

Procedure

Monitor your replication using the CLI.

stats --replication --partner *partner_name*

Example

stats --replication --partner rep_array5

Syslog

Syslog is a standard for computer message logging. It is supported on a variety of devices and platforms, and is used to store management, security, informational, debugging, and other types of messages about these devices.

The syslog stores important information such as records of administrator manipulation of the storage array, and a history of alerts or issues with the array. Using syslog, system log files can be shipped from an array group to a centralized, remote server. The benefits of this include:

- Cost savings - system log files can be archived on inexpensive media rather than on the array.
- Ease of use - a central repository consolidates data from multiple arrays into one area, so it is not necessary to log into every array to get the data.
- Data analytics - it's easier to examine logs for troubleshooting, security, and health-related issues if they are on a central device.

With syslog enabled, arrays can communicate with third party monitoring tools without the need of custom code because it uses the standard syslog protocol.

Arrays support the Red Hat Enterprise Server and Splunk implementations of syslog. UDP is used to communicate between the array group and the syslog server (SSL is not supported at this time). One syslog message is generated for each alert and audit log message. Alert severity types include INFO, WARN and ERROR.

Enable Syslog

Note: To enable syslog you must have Power User privileges or higher.

The command to enable syslog is a suboption of the **group --edit** option.

Procedure

1. Log into the array.
2. Enable syslog.

   group--edit--syslog_enabledyes

   Syslog is now enabled for this array.

3. Specify the syslog server.

   group--edit--syslog_server*syslog_server*

   where *syslog_server*is a valid hostname or IP address of the syslog server you will use.

4. Specify the syslog port.

   group--edit--syslog_port*syslog_port*

   where *syslog_port* is a valid integer *0-65535*

   No check is performed to determine whether the host name exists or the IP address exists or is valid, or whether the port number is valid.

Disable Syslog

Note: To disable syslog you must have Power User privileges or higher.

The command to disable syslog is a suboption of the **group --edit** option.

Procedure

1. Log into the array.
2. Disable syslog.

   group--edit--syslog_enabledno

Audit Log Management

The audit log keeps records of all user-initiated non-read operations performed on the array, and which user performed the operation. You can search the audit log by activity and object type, name or both. You can also filter the audit log by time range, username, activity category, and access type. Administrators can view the audit log in a summary table with faceted browsing by time, activity category, and across access type.

Audit logging has changed from version 2.2.3.0, including which operations are audited, and syslog message format. Operations are not audited on non-group leader arrays, or on the standby controller of the group leader array, to which only the root user has access. In addition, console logout is not audited. Operations cannot be logged before the group is set up, which is when audit logging begins.

Audit logs, along with alerts, are posted to a syslog server if one is configured, using the following format:

Jan 22 17:51:01 sjc-b11-va-B NMBL: Group:group-sjc-b11-va Type:2001 Time:Thu Jan 22 17:51:01 2015#012 Id:275 Object Id:- Object:vol-10 Access Type:pam Client IP:10.20.20.248 Status:Succeeded

Audit log messages are not sent through emails, SNMP traps, or to InfoSight in real time. However, error messages for failed operations are converted to HTTP-like errors.

Audit logs are merged during a group merge, beginning with the users. Users from the source group are remapped to new users in the destination group. After the users are merged, the audit logs are merged.

The audit log is automatically purged. When the count reaches 21,000, an alert is sent warning that a purge will occur when the count reaches 24,000. At 24,000 messages, the oldest 5000 messages are purged (the most recent 19,000 log entries are kept).

Audit Log Panel

Users with the Administrator role can access the audit log page by selecting **Monitor** > **Audit Log**.

The main audit log page has two panels - a summary table panel on the right that provides a list of audit log records, and a collapsible facets panel on the left used to narrow down audit log records in the table. When the panel is collapsed, any facet settings remain in effect.

Facets Panel

The facets panel provides four ways to filter content:

- Search by activity or object - Enter words (case insensitive) to search by activity or object. The log table list changes depending on the words you enter. You will not be able to search on deleted users, root users, or system users.
- Date Range - Select from a dropdown list of common time intervals (All, Last Hour, Last 24 Hours, Last 7 Days, Last 30 Days, and Custom...). The default value is *All*. *Last* means last from the current time. Selecting the Start Time and End Time fields under Custom Displays allows you to specify a date and time range. Any values you enter remain in effect when *All* or *Last* is selected. Audit log records are filtered whenever you make a selection from the dropdown list or enter a valid start and end date after selecting Custom.
- Activity Category - Check up to six checkboxes (Data Provisioning, Data Protection, Data Access, User Access, System Configuration, Software Update) to filter the log table list by type of activity audited. Audit log records are filtered when a checkbox is checked or unchecked.
- Access Type - Check up to three checkboxes (API, CLI, GUI) to filter the log table list by type of access audited. Audit log records are filtered when a checkbox is checked or unchecked.

Summary Table

The summary table has six columns which can be used to sort or filter the data:

- Time - Provides sorting and filtering of when activities take place. The Time filter is the same as the one provided in the facets panel.
- Activity - Provides filtering of user actions. You can use the Activity filter to further refine what you have selected in the facets panel.
- Status - Provides sorting of operation status icons (successful, in process, or failed). Hovering the mouse over a failed status icon displays a tooltip describing the cause of the failure.
- User - Provides sorting of full names of registered users. Clicking the hyperlinked username brings up the user details page.
- Client IP Address - Provides sorting of the IP addresses where the activity was invoked.
- Access Type - Provides sorting and filtering of access methods (API, CLI, GUI). The Access Type filter is the same as the on provided in the facets panel.

The summary table is refreshed whenever you change the facet panel settings. You can also refresh the table by clicking the refresh icon in the upper right corner of the panel. The table does not refresh automatically.

Note: System users (such as VSS agent) are shown as *<system>* in the username column and *System* in the User Full Name column. User information can be empty if the authentication failed (for example, from an expired session).

User Management

Users with the Administrator role can access the Manage Users page by selecting **Administration** > **Security** > **Manage Users**. This page shows an audit log summary table for the selected user. All user's activities are displayed. To show new user activity, reload the table or reselect the user.

For other tasks you can perform from this page, see Role-Based Access Control on page 116.

Disaster Recovery

For disaster recovery, you must have at least one additional Nimble array that is configured as a replication partner. Replication partners can serve data to the original initiators while the original array is inaccessible. By including multiple Nimble arrays in your network, you can quickly restore access to data even in case of a catastrophic failure.

In the unusual event of a complete failure of the array, set the replication partner online and point your initiators to the volumes on the replication partner. Your volume data is available to applications during the recovery of the data to the failed array.

The two methods to move operations from one volume collection to its replication partner are handover and promotion/demotion. If the original array is accessible, handover is always preferred.

In case of a complete failure, within the Site Recovery Manager (SRM) the "failback" procedure is to delete the failed VM from vCenter, then replicate the backed-up LUNs to a new VM.

Important: Follow the procedures in this section to perform disaster recovery using the Nimble array only. For information on disaster recovery in HPE Cloud Volumes, refer to the HPE CV portal documentation.

Handover Overview

A *handover* is a controlled way to migrate all volumes associated with a particular volume collection to a replication partner without any loss of data. The new owner must be an immediate downstream replication partner for replicas based on this volume collection.

Note: To use handover, both replication partners must be active and functional.

A handover instructs a downstream replication partner to become the upstream replication partner and provide the iSCSI initiators access to volumes. It allows the replication partner to take ownership of a volume collection.

The results of handing over a volume collection are:

- The volumes associated with the volume collection on this array are set offline.
- Snapshots of the associated volumes are taken.
- The snapshots are replicated to the downstream replication partner.
- Volume collection ownership is transferred to the replication partner.
- The volumes associated with the volume collection on the replication partner are set online.

By default, the direction of replication is automatically reversed. That is, when the downstream group becomes the owner of the volume collection, the upstream group becomes its replication partner.

Perform a Handover

Procedure

1. Perform a handover.

   volcoll --handover *name of volume collection* **--partner** *group name*

   volcoll --handover VolumeCollection5 --partner group26

   Note: If you do not want to reverse the direction of replication, add the --no_reverse option.

   The message is displayed: Handover is in progress.

   The new owner is the replication partner group.

2. At the command prompt, type:

   volcoll --info *name of volume collection*

   volcoll --info VolumeCollection5

3. Under Schedule, verify that the group names listed as Owned by: and Replicate to: have reversed.

What to do next

If the replication partners use different access protocols, additional steps are needed to bring the replica online. For example, configuring multi-initiator access is applicable for iSCSI arrays, but is inapplicable for Fibre Channel arrays. By default, Fibre Channel allows multi-initiator access, and you cannot disable it. Access control is done through LUN mapping.

Make a Replica Available to Applications

Before you begin

The original volume must be offline, whether due to disaster or as part of a planned outage. If this is a planned outage, perform a replication immediately prior to failing the original array.

Procedure

Take the volume offline.

vol --offline *volname*

Results

The initiators now use the volume on the remote array to read and write data. After the original array is restored, use handover to hand control back to the original.

Promote a volume collection

Important: You should only use promotion when the original array is unavailable.

Note: When promoting a volume collection to a replicated volume, the ACLs are copied to the promoted set.

Procedure

Initiate a promotion.

volcoll --promote *volcollname*

Results

To replicate the volume collection back to the original primary from the newly promoted array, demote the volume collection on the original array. By default, the replication direction is reversed.

Demote a volume collection

Procedure

Initiate a demotion.

volcoll --demote *volcollname*

Claim a volume

Claiming a volume lets you take ownership of a formerly replicated volume on the downstream replication partner that is no longer part of a volume collection. Without claiming the volume, you cannot make any changes to the volume attributes.

Claim should be used if a primary upstream partner is no longer present and access to the replicated volume is required at the downstream site, and or if you want to migrate this volume replica to a new volume collection belonging to the downstream partner system.

Procedure

Claim the volume.

vol --claim *vol_name*

Array Administration

Array administration involves performing many array-based administrative tasks. Examples include email alert configuration, password management, timeout activity, SNMP, and HTTP proxy settings.

Configure Email Alerts

You can configure email alerts to use regular or secure Simple Mail Transfer Protocol (SMTP) processing, depending on which mode is appropriate for your environment.

You might choose to use regular SMTP for email alerts if you have an SMTP server installed on your network that accepts email messages from external parties. You might choose to use secure SMTP if you have an SMTP server installed on your network, but prefer to disallow anonymous relays, or if you do not have an internal email server because you implemented cloud-based email, such as Office 365.

You can configure email alerts differently for each group of arrays. You must have at least Power User permission to configure SMTP-based email alerts.

Using the following command options you can modify group settings, change the hostname or IP address of the SMTP server or its port number, allow use of SMTP authentication and encryption (with options to change the authentication password and enable the level of encryption, when enabled), change the email address for sending and user receiving email alerts, enable email to be sent to support, and set the minimum event data alert level at which alerts are sent.

For a full list of all group --edit options, refer to the *Command Reference*.

Procedure

Configure email alerts.

group--edit [**--smtp_server** *smtp server*] [**--smtp_port** *smtp port*] [**--smtp_auth** {**yes** | **no**}] [**--smtp_username** *username*] [**--smtp_encrypt_type** {**none** | **starttls** | **ssl**}] [**--smtp_from_addr** *email addr*] [**--smtp_to_addr** *email addr*] [**--send_event_data** {**yes** | **no**}] [**--alert_level** {**info** | **warning** | **critical**}]

Note: When you edit SMTP, you must always include the --smtp_username, --smtp_encrypt_type, and password to prevent the operation from failing. If you are using authorization, you must also include --smtp_auth yes. You do not enter the password on the command line. After you run the command, you will be asked to enter the password. Providing the password this way avoids having to display it in plain text.

Note: The syntax for configuring SMTP email to send alert messages for critical alerts to multiple email addresses is, for example:

```
--smtp_to_addr Fred@here.com --smtp_to_addr George@here.com
```

Example

Configuring SMTP email to send alert messages for critical alerts from user Joe Black to support.

```
$ group --edit --smtp_auth yes --smtp_username JoeBlack@mycompany.com --
smtp_encrypt_type ssl --smtp_from_addr Joe.Black@customer.com --smtp_to_addr
support@nimblestorage.com --alert_level critical
```

Diagnostics for Nimble Analytics

Diagnostics for Nimble Analytics (DNA) collects product operational data including the performance, reliability, and configuration characteristics of the array and sends this information to Nimble Storage once per day. The information is used for proactive monitoring, analysis, and problem resolution. No user data is ever accessed or collected by DNA.

By default, DNA is enabled. Leaving DNA enabled is strongly recommended because this allows Nimble Storage Support personnel to continually monitor the health of the array and recommend corrective actions in case of any issues.

Enable Autosupport

Enable autosupport.

Procedure

group --edit [**--autosupport yes**]

Disable Autosupport

If you disable autosupport, no statistics or diagnostics are sent to support. Leaving autosupport enabled is strongly recommended.

Procedure

group --edit [**--autosupport no**]

Manually Send an Autosupport

If you are asked to do so, you can send an autosupport at any time for analysis and problem resolution. No user data is ever accessed or collected by autosupport.

Procedure

group --autosupport_initiate

Enable a Secure Tunnel

Support may request a tunnel into the array to help troubleshoot an issue. You will be able to enable a secure tunnel. By default, secure tunnels are disabled. Once enabled, support can open a tunnel.

Procedure

group --edit --support_tunnel yes

Configure a Proxy Server

Some features require a secure Internet connection. If your network requires the use of a general proxy server or an HTTP proxy server for secure connections, configure the array to use the correct server.

Procedure

group --edit --proxyserver *server_name* **--proxyport** *port* **--proxuser** *username* **--proxpassword** *password*

Example

```
$ group --edit --proxyserver proxy.hpe.com --proxyport 8080
```

Change an Array Name

Changing the array name does not change the array serial number or any other information.

Procedure

array --edit [**--name** *array_name*]

Set Up SNMP

The array uses SNMP to communicate with network management systems. It supports SNMP versions 1, 2, and 2c. However, the device sends traps but does not receive them. You can download the SNMP MIB from the support site.

Note: The array uses the alert level settings for email alerts to determine the events that are sent as SNMP traps.

SNMP sends information to the network in one of two ways. Using the first method, the network management system sends a request to retrieve information and then subsequently receives a response. Using the second method, the traps are sent automatically, based on trap level settings.

Procedure

1. Enable SNMP traps.

   group --snmp_trap_host

2. Configure SNMP trap destination.

   array --snmp_trap_dest

Fail Over a Controller

A failover switches management of the array from the active controller to the standby controller.

While you can manually perform a failover, a failover can also be system driven. For example, in an iSCSI array, a failover can occur when the standby controller has better connectivity. In a Fibre Channel array, a failover will occur when the active controller loses all connectivity.

Note: A failover will not start automatically until connectivity to the array is lost for ~6 seconds.

You must perform a failover during a controller upgrade or when directed by support.

Before you begin

Failover requires one controller to be in Active mode and the other controller to be in Standby mode.

Procedure

1. Determine the name of the array.

   array --list

   The name of the array appears in the list.

2. Perform the failover.

   failover --array *array_name*

   During the failover operation, the standby controller first goes into Solo mode, and then into Active mode. The active controller goes into Unknown mode, then into Stale mode, and finally into Standby mode.

Example

```
$ failover --array datamaster-a
```

Shut Down an Array

You can gracefully shut down your array. If you shutdown an array used in peer persistence, manually hand over each volume before shutting down the array. If you do not, all volumes become unavailable.

Procedure

Shut down the array.

halt --array

Note: When shutting down an array, make sure you manually power off any expansion shelves.

Alarm Management

Alarms help users to better monitor and manage their storage by alerting them to a variety of different events. Compared to events, which are presented in a log based on the sequence of the occurrence of the event, alarms are active issues on the array presented in real time.

The alarms have multiple states including open, acknowledged, and closed. You can perform the following actions with the alarms:

- Set and modify reminders for the alarms
- Mute reminders for a period of time
- Set or modify the frequency of the reminders
- Disable the reminders

The alarms have numerous properties, including the following:

- Event time
- Severity
- Category
- Description
- Object type
- Object name
- Status
- Username of the acknowledging user
- Array name
- Group name
- Time of acknowledgment
- Recovery time
- Object ID

  Object IDs might not exist for events that are raised by platform.

List Alarms

You can use the **alarm** command with the --list option to list all of the alarms on the array. You can also use the **alarm** --list command to find out the ID of a particular alarm and use that ID to acknowledge, modify or delete an alarm.

Procedure

From a command prompt, request a list of alarms on the array.

alarm --list

The command returns a list of all alarms on the array, their ID, severity, time, status, and details.

```
$ alarm --list
------------+--------------+-------------------------------+-----------------
-------+-------------------+--------------------------------
    ID           Severity       Time                                  Status
              Array                 Detail
------------+--------------+-------------------------------+-----------------
-------+-------------------+--------------------------------
       38        CRITICAL  Oct 10 2016 16:49:30     Open
```

```
AF-110068       Volume test string from alert gen space usage is at 15%
and approaching quota of 20%. It will be taken offline if it exceeds the
quota

       39     WARNING   Oct 13 2016 17:32:01     Acknowledged   AF-110068
     IP interface 10.19.0.110 down on controller B NIC port eth1

       40     WARNING   Oct 19 2016 15:42:32     Acknowledged   AF-110068
     IP interface 172.19.0.137 down on controller A NIC port eth6

       41     WARNING   Nov  3 2016 16:33:45     Open                   AF-
110068         IP interface 1.1.1.1 down on controller B NIC port eth1.11

      42      WARNING   Nov  3 2016 16:35:41     Open                   AF-
110068         IP interface 1.1.1.7 down on controller B NIC port eth1.11
```

Enable and Disable the Alarm Feature

Note: You can disable the alarm feature in the CLI only. You cannot disable the alarm feature in the GUI.

You must use the **group** command to enable or disable the alarm feature.

Procedure

From a command prompt, enable or disable the alarm feature.

Using the **group** command, you can either enable or disable an alarm feature.

group --edit --alarms_enabled *yes|no*

Acknowledge Alarms

You can set the frequency of alarm notifications at the time that you acknowledge the alarm.

Procedure

From a command prompt, acknowledge a particular alarm.

You can also include sub-commands to set the frequency of the notifications.

alarm --acknowledge *ID* --remind_every *period* --remind_every_unit {*minutes|hours|days*}

Example

The following example lists all the alarms on the array (there is only one) and then acknowledges that alarm setting the notifications of the alarm to remind the user every day.

```
$ alarm --list
------------+-------------+-------------------------------+------------+-----
-------------+---------------------------------
ID             Severity       Time                              Status
Array                Detail
------------+-------------+-------------------------------+------------+-----
-------------+---------------------------------
       5        CRITICAL  Oct 20 2016 16:49:30     Open         -
          Volume test string from alert gen space usage is at 0%
and approaching quota of 0%. It will be taken offline if it exceeds the quota
```

```
$ alarm --acknowledge 5 --remind_every 1 --remind_every_unit days
```

Unacknowledge Alarms

You can use the **– unacknowledge** command if you previously acknowledged an alarm.

Procedure

From the command line prompt, unacknowledge a particular alarm.

alarm --unacknowledge *ID*

The CLI returns output only if there is an error. To verify that the alarm has been successfully unacknowledged, you can run the **alarm** --info command and check the status of the alarm.

Example

The following example shows the original list of alarms. Alarm ID 6 is in the Acknowledged status. The following command then unacknowledges alarm ID 6, which is shown in the new list of alarms on the array. The new status of alarm ID 6 is now Open.

```
$ alarm --list
----------+-----------------+------------------------------+-----------------
----------+----------+-------------------------------------------
ID           Severity           Time                                Status
                    Array      Detail
----------+-----------------+------------------------------+-----------------
----------+----------+-------------------------------------------
       6      CRITICAL      Oct 31 2016 17:22:17     Acknowledged
-          Volume test string from alert gen
space usage at 0% is over the configured quota at 0%

       8      WARNING    Oct 31 2016 17:23:26     Open
     -          The amount of used space in the
pool test string from alert gen has reached 0% of free space.

       9      WARNING     Nov  1 2016 10:49:02     Open
      -          Configuration synchronization to
array test string from alert gen delayed, will continue to retry

$ alarm --unacknowledge 6

$ alarm --list
----------+-----------------+------------------------------+-----------------
----------+----------+-------------------------------------------
ID           Severity           Time                                Status
                    Array      Detail
----------+-----------------+------------------------------+-----------------
----------+----------+-------------------------------------------
       6      CRITICAL      Oct 31 2016 17:22:17     Open
         -          Volume test string from alert gen
space usage at 0% is over the configured quota at 0%

       8      WARNING    Oct 31 2016 17:23:26      Open
      -          The amount of used space in the
pool test string from alert gen has reached 0% of free space.

       9      WARNING     Nov  1 2016 10:49:02     Open
```

```
      -          Configuration synchronization to
array test string from alert gen delayed, will continue to retry
```

Change an Alarm Reminder

You use the **alarm** –edit command to change the frequency of alarm notifications. The alarm must have been acknowledged before you can modify the reminder. If an alarm has not already been acknowledged, you can set the notification frequency when you acknowledge the alarm.

Procedure

1. From the command prompt, edit the alarm notification frequency.

   alarm --edit *id* **--remind_every** *period* **--remind_every_unit** *minutes|hour|days*

   There is no output or acknowledgement of this change.

2. View the current notifications for a particular alarm.

   alarm --info *id*

Example

The following example shows the command to modify the notification of alarm ID 9 to be tripped every four hours.

$ alarm --edit 9 --remind_every 4 --remind_every_unit hours

Delete Alarms

Note: Be careful using this command because the alarms indicate severe conditions on the system.

Procedure

Delete the alarm specified by the ID associated with the alarm.

alarm --delete *ID*

To verify that the alarm was deleted, run the **alarm** --list command.

Events

The array monitors events and displays them on the Events page. Events can let you know when something needs your attention, or when an event may be about to occur. They are an excellent diagnostic aid when you attempt to locate the source of a problem or potential problem on the array.

The array provides two locations from which you can view events: the events summary and recent events, as shown on the Home page, and a list of all events that you can filter as shown on the *Events* details page. Each event has a priority that you can use to filter information in the list, as well as determine whether or not the event requires manual intervention.

Event Severity Levels

The array uses the standard alert system that supports three basic levels of severity. Depending on the level of the event that has been logged, immediate action could be required.

Table 13: Event Severity Levels

Severity Level	Description	Examples
All	All events are shown, regardless of severity. Manual intervention may or may not be needed.	
Critical	An event has occurred that requires immediate attention. Data loss or hardware damage may occur if action is not taken quickly. Critical alerts also trigger email notification, defined on the Administration tab.	The system has reached space capacity. A drive has failed.
Warning	An event has occurred that might impact system performance. Action will likely be necessary, but no damage will occur if the action is not taken immediately.	A scheduled snapshot was not completed successfully. A drive is experiencing write errors.
Info	An event has occurred that does not require action to be taken and does not affect system performance. This level of event is useful for troubleshooting or for determining system trends.	The administrator password was changed. A controller was restarted.

View Events

You can use the **alert** command to view a list of events that occurred on the array. You can also use the --severity option to filter the list based on the alert level.

Procedure

List events using the CLI.

alert --list --severity {*info* | *warning* | *critical*}

The --severity option is not required.

Events and Alert Messages

Table 14: Alert Configuration

Configuration			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10244	Warning	Group	"Configuration synchronization to array delayed will continue to retry"
10245	Information	Group	"Configuration synchronization completed for array" "Configuration synchronization completed on array %s"
10246	Warning	Controller	"Group Leader Platform Configuration synchronization is delayed will continue to retry"
10247	Information	Controller	"Group Leader Configuration synchronization completed" "Group Leader Configuration synchronization completed on active controller %s"
10248	Warning	Controller	"Configuration synchronization to array delayed will continue to retry"
10249	Information	Controller	"Configuration synchronization completed for array" "Configuration synchronization completed on array %s"
10267	Critical	Group	"Encryption deactivated" "Encryption deactivated. Encrypted volumes cannot be accessed or created. Enter encryption passphrase to reactivate."
10268	Information	Group	"Encryption master key was deleted" "Encryption master key was deleted. No data was lost because no encrypted volumes exist."
10269	Warning	Group	"Encryption mode was changed to secure mode" "Encryption mode was changed to secure mode. An array reboot will require passphrase entry."
10270	Warning	Group	"Encryption mode was changed to available mode" "Encryption mode was changed to available mode. An array reboot will not require passphrase entry."

Configuration			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10271	Information	Group	"Encryption master key was created" "Encryption master key was created. Encrypted volumes can now be created."
10272	Information	Group	"Encryption passphrase was changed" "Encryption passphrase was changed."
10273	Information	Group	"Encryption cipher was changed" "Encryption default cipher was changed to %s."
10274	Information	Group	"Encryption scope was changed" "Encryption default scope was changed to %s."
10275	Information	Group	"Encryption activated" "Encryption activated. Encrypted volumes can now be accessed and created."
10276	Warning	Group	"Encrypted volume access may be slow" "Array %s is a %s model. Encryption on this array may be slow."
10277	Critical	Group	"Encryption master key was deleted" "Encryption master key was deleted. Encrypted volumes are now permanently inaccessible."
10280	Warning	Initiator Group	"Initiator group synchronization is delayed to arrays in the group" "Initiator group %s synchronization is delayed to arrays in the group. ""Synchronization will be retried. Additional edits to initiator groups and creation of new access control records will fail during the delay. ""When synchronization succeeds
10281	Information	Initiator Group	"Initiator group synchronization to arrays in the group succeeded" "Initiator group %s synchronization to arrays in the group succeeded."
10500	Warning	Group	"Object count reached maximum limit" "Number of %s in %s has reached the maximum limit of %ld"
10501	Information	Group	"Object count under maximum limit" "Number of %s in %s is now under the maximum limit of %ld"

Configuration			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10502	Warning	Group	"Object count over warning limit" "Number of %s in %s has reached the warning limit of %ld"
10503	Information	Group	"Object count under warning limit" "Number of %s in %s is now under the warning limit of %ld"
10504	Warning	Group	"Object count reached maximum limit." "Number of %s %s in %s has reached the maximum limit of %s."
10505	Information	Group	"Object count under maximum limit." "Number of %s %s in %s is now under the maximum limit of %s."
10506	Warning	Group	"Object count over warning threshold." "Number of %s %s in %s has reached the warning threshold of %s."
10507	Information	Group	"Object count under warning threshold." "Number of %s %s in %s is now under the warning threshold limit of %s."

Table 15: Alert Hardware

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
2022	Warning	Controller	"Attempting controller failover" "Attempting failover to active role because controller %s has %s."
2023	Warning	Controller	"Controller could not complete failover" "Attempted failover did not succeed. Check the %s on controller %s."
2024	Information	Controller	"Controller failover occurred" "Failover to active role occurred. Controller %s is now the active controller. Check %s on controller %s."
2028	Warning	Controller	"Controller rebooted unexpectedly" "Controller %s rebooted unexpectedly. Contact Nimble Storage Support."
12000	Critical	Temperature	"Overtemperature shutdown" "Overtemperature on controller % s

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12001	Critical	Temperature	"Controller overtemperature" "Overtemperature on controller %s (%d Celsius)"
12002	Critical	Temperature	"Backplane over-temperature" "Overtemperature on backplane (%d Celsius)"
12003	Warning	Array	"Flash memory in head shelf does not meet the minimum requirement" "Amount of flash memory detected on head shelf does not meet the minimum requirement ""for this array model. Performance may be degraded. Check the SSDs for incompatible sizes or failures."
12004	Information	Array	"Flash memory in head shelf meets the minimum requirement" "Amount of flash memory detected on head shelf now meets the minimum requirement ""for this array model."
12100	Critical	Disk	"Disk failed" "Disk %s failed at slot % d
12101	Warning	Disk	"Disk failed" "Disk %s failed at slot %d"
12102	Warning	Disk	"Disk missing" "Disk %s missing at slot %d"
12103	Information	Disk	"Disk added" "Disk %s added at slot %d"
12104	Information	Disk	"Disk removed" "Disk %s at slot %d is removed by user using cli"
12105	Warning	Disk	"SSD failed" "SSD %s failed at slot %d"
12106	Warning	Disk	"SSD missing" "SSD %s missing at slot %d"
12107	Information	Disk	"SSD added" "SSD %s added at slot %d"
12108	Information	Disk	"SSD removed" "SSD %s at slot %d is removed by user using cli"
12109	Warning	Disk	"Foreign disk detected" "Disk %s is not being used

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12110	Warning	Disk	"Foreign SSD detected" "SSD %s is not being used
12111	Warning	Disk	"Invalid disk size for this model" "Disk %s of size %d GB is ""not the right size for this model
12112	Warning	Disk	"Disk failed" "Disk %s failed on %s shelf %s at slot %d"
12113	Warning	Disk	"Disk missing" "Disk %s missing"
12114	Information	Disk	"Disk added" "Disk %s added on %s shelf %s at slot %d"
12115	Information	Disk	"Disk removed" "Disk %s on %s shelf %s at slot %d is removed by ""user using cli"
12116	Warning	Disk	"SSD failed" "SSD %s failed on %s shelf %s at slot %d"
12117	Warning	Disk	"SSD missing" "SSD %s missing"
12118	Information	Disk	"SSD added" "SSD %s added on %s shelf %s at slot %d"
12119	Information	Disk	"SSD removed" "SSD %s on %s shelf %s at slot %d is removed by ""user-issued CLI command"
12120	Critical	Disk	"Only one SSD drive left" "Only a single SSD drive is left in the system"
12121	Warning	Disk	"Disk failed" "Disk %s of capacity %dTB failed on %s shelf %s at slot %d."
12122	Information	Disk	"Disk added" "Disk %s of capacity %dTB added on %s shelf %s at slot %d."
12123	Information	Disk	"Disk removed" "Disk %s of capacity %dTB on %s shelf %s at slot %d was removed by ""user-issued CLI command."

Hardware

Alert/Event	Severity	Type	Description/Possible Cause/Actions
12124	Warning	Disk	"SSD failed" "SSD %s of capacity %dGB failed on %s shelf %s at slot %d."
12125	Information	Disk	"SSD added" "SSD %s of capacity %dGB added on %s shelf %s at slot %d."
12126	Information	Disk	"SSD removed" "SSD %s of capacity %dGB on %s shelf %s at slot %d was removed by ""user-issued CLI command."
12127	Warning	Disk	"Invalid disk size for this shelf" "Disk %s of size %d GB is the wrong size on %s shelf %s at slot %d. ""This shelf needs a %d GB disk."
12128	Warning	Disk	"Invalid SSD size for this shelf" "SSD %s of size %d GB is the wrong size on %s shelf %s at slot %d. ""Replace SSD with supported model or contact Nimble Storage Support."
12129	Warning	Disk	"Disk not supported on All Flash Shelf" "Disk %s of size %d TB is not supported on All Flash Shelf %s at slot %d."
12130	Warning	Disk	"Disk not supported on All-Flash Shelf" "Disk %s of size %d TB is not supported on All-Flash Shelf %s at slot %d."
12131	Critical	Disk	"No SSDs found" "No SSDs found. This can cause performance degradation and possibly ""service outages. Insert functioning SSDs in the array and verify that the array recognizes ""the drives."
12132	Warning	Disk	"Disk(s) inaccessible" "One or more disks are inaccessible from controller %s. ""The disks may be missing or not responding. Failover and software ""updates may be affected. Contact Nimble Storage Support."
12133	Warning	Disk	"Disk taken out of service for diagnosis" "%s %s of capacity %s on %s shelf %s at slot %d ""is experiencing errors and is being temporarily taken out of service for inspection ""and attempted recovery. No action required at this time."

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12134	Warning	Disk	"SSD failed" "SSD %s of capacity %dGB failed on %s shelf %s at slot %s."
12135	Information	Disk	"SSD added" "SSD %s of capacity %dGB added on %s shelf %s at slot %s."
12136	Information	Disk	"SSD removed" "SSD %s of capacity %dGB on %s shelf %s at slot %s was removed by ""user-issued CLI command."
12137	Warning	Disk	"Invalid disk size for this shelf" "Disk %s of size %d GB is the wrong size on %s shelf %s at slot %s. ""This shelf needs a %d GB disk."
12138	Warning	Disk	"Invalid SSD size for this shelf" "SSD %s of size %d GB is the wrong size on %s shelf %s at slot %s. ""Replace SSD with supported model or contact Nimble Storage Support."
12139	Warning	Disk	"Disk not supported on All-Flash Shelf" "Disk %s of size %d TB is not supported on All-Flash Shelf %s at slot %s."
12140	Warning	Disk	"Disk taken out of service for diagnosis" "%s %s of capacity %s on %s shelf %s at slot %s ""is experiencing errors and is being temporarily taken out of service for inspection ""and attempted recovery. No action required at this time."
12200	Warning	NIC	"IP interface down" "IP interface %s down on controller %s NIC port %s"
12201	Information	NIC	"IP interface up" "IP interface %s up on controller %s NIC port %s"
12202	Warning	NIC	"Group IP interface unavailable" "Group IP interface unavailable"
12203	Warning	NIC	"All Data IP interfaces unavailable" "All Data IP interfaces unavailable"
12204	Critical	NIC	"IP connectivity lost all links are down"

Hardware

Alert/Event	Severity	Type	Description/Possible Cause/Actions
12205	Warning	NIC	"Migrating subnet NIC configuration" "NIC %s for subnet %s does not exist. Migrating configuration to NIC %s. Edit configuration if necessary."
12206	Warning	NIC	"NIC for subnet missing" "NIC %s for subnet %s does not exist. Edit configuration to ensure subnet is configured for an existing NIC."
12207	Warning	NIC	"Migrating IP NIC configuration" "NIC %s for IP %s does not exist. Migrating configuration to NIC %s. Edit configuration if necessary."
12208	Warning	NIC	"NIC for IP missing" "NIC %s for IP %s does not exist. Edit configuration to ensure IP is configured for an existing NIC."
12209	Warning	NIC	"Migrating route NIC configuration" "NIC %s for route %s gateway %s does not exist. Migrating configuration to NIC %s. Edit configuration if necessary."
12210	Warning	NIC	"NIC for route missing" "NIC %s for route %s gateway %s does not exist. Edit configuration to ensure route is configured for an existing NIC."
12211	Warning	NIC	"Standby controller has better network connectivity. Failing services over to standby controller" "Standby controller has better network connectivity. Failing services over to standby controller"
12212	Warning	NIC	"Duplicate IP Address Detected" "A duplicate of IP address %s is found in the network."
12213	Warning	NIC	"Discovery IP interface unavailable" "Discovery IP interface unavailable"
12214	Warning	NIC	"Target IP interfaces unavailable" "Target IP interfaces unavailable"
12215	Warning	NIC	"Link down for Group IP interface(s)" "Link down for Group IP interface(s)"
12216	Warning	NIC	"Link down for Discovery IP interface(s)" "Link down for Discovery IP interface(s)"

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12217	Warning	NIC	"Link down for Target IP interface(s)" "Link down for Target IP interface(s)"
12218	Warning	NIC	"Link down for Data IP interface(s)" "Link down for Data IP interface(s)"
12219	Warning	NIC	"Link down for iSCSI Data IP interface(s)" "Link down for iSCSI Data IP interface(s)"
12220	Warning	NIC	"Link down for cluster Data IP interface(s)" "Link down for cluster Data IP interface(s)"
12221	Warning	NIC	"Unresponsive NIC Port Detected" "NIC Port %s is unresponsive."
12222	Critical	NIC	"IP connectivity lost all links down on a subnet"
12223	Critical	NIC	"Network connectivity lost all links down on a subnet"
12224	Warning	NIC	"NIC migration failed" "Migration of NIC port(s) %s failed. This could likely happen if array hardware updates resulted in reduced number of NICs."
12301	Warning	Temperature	"Sensor: Alert raised" "%s : Min %d Max %d (over %d hours %d mins) Last Reading %d"
12302	Information	Temperature	"Sensor: Alert cleared" "%s : Min %d Max %d (over %d hours %d mins) Last Reading %d"
12303	Warning	Temperature	"Sensor: Alert raised" "%s : failure "
12304	Information	Temperature	"Sensor: Alert cleared" "%s : failure cleared "
12305	Warning	Temperature	"Sensor: is missing?: " "%s"
12306	Warning	NVRAM	"NVRAM battery is disabled" "NVRAM Battery %d is disabled on controller %s"

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12307	Critical	NVRAM	"NVRAM batteries are disabled" "NVRAM %d batteries are disabled on controller %s"
12308	Information	NVRAM	"NVRAM battery is OK" "NVRAM Battery %d is OK on controller %s"
12310	Warning	Temperature	"System temperature is warm" "current temperature for ""sensor at %s is %d C"
12311	Warning	Temperature	"System temperature is cold" "current temperature for ""sensor at %s is %d C"
12312	Information	Temperature	"System temperature is ok" "System temperature is now ""within range: current temperature for sensor at %s is %d C"
12313	Warning	Temperature	"Temperature sensor failed" "Temperature sensor at %s ""is not operational"
12314	Information	Temperature	"Temperature sensor is ok" "Temperature sensor at %s is ""operational: current temperature is %d C"
12315	Information	Cooling Fan	"Fan in high speed" "Fan at %s %s is running fast: ""current speed %d rpm"
12316	Information	Cooling Fan	"Fan in low speed" "Fan at %s %s is running slow: ""current speed %d rpm"
12317	Information	Cooling Fan	"Fan is ok" "Fan at %s %s is now running at speed ""%d rpm"
12318	Warning	Cooling Fan	"Fan stopped" "Fan at %s %s stopped working"
12319	Warning	Cooling Fan	"Fan missing" "Fan at %s %s is missing"
12320	Warning	Power Supply	"Power supply fail" "Power supply at %s is either not ""connected or failed"
12321	Warning	Power Supply	"Power supply missing" "Power supply at %s is missing"

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12322	Information	Power Supply	"Power supply is ok" "Power supply at %s is ok"
12323	Warning	Temperature	"System temperature is warm" "current temperature for ""sensor \"%s\" at %s is %d C"
12324	Warning	Temperature	"System temperature is cold" "current temperature for ""sensor \"%s\" at %s is %d C"
12325	Information	Temperature	"System temperature is ok" "System temperature is now ""within range: current temperature for sensor \"%s\" at %s is ""%d C"
12326	Warning	Temperature	"Temperature sensor failed" "Temperature sensor \"%s\" ""at %s is not operational"
12327	Information	Temperature	"Temperature sensor is ok" "Temperature sensor \"%s\" ""at %s is operational: current temperature is %d C"
12328	Information	Cooling Fan	"Fan in high speed" "Fan \"%s\" at %s %s is running ""fast: current speed %d rpm"
12329	Information	Cooling Fan	"Fan in low speed" "Fan \"%s\" at %s %s is running ""slow: current speed %d rpm"
12330	Information	Cooling Fan	"Fan is ok" "Fan \"%s\" at %s %s is now running ""at speed %d rpm"
12331	Warning	Cooling Fan	"Fan stopped" "Fan \"%s\" at %s %s stopped working"
12332	Warning	Cooling Fan	"Fan missing" "Fan \"%s\" at %s %s is missing"
12333	Warning	Power Supply	"Power supply fail" "Power supply \"%s\" at %s is ""either not connected or failed"
12334	Warning	Power Supply	"Power supply missing" "Power supply \"%s\" at %s is ""missing"

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12335	Information	Power Supply	"Power supply is ok" "Power supply \"%s\" at %s is ok"
12336	Warning	Temperature	"System temperature is warm" "current temperature for ""sensor \"%s\" on shelf \"%s\" at %s is %d C"
12337	Warning	Temperature	"System temperature is cold" "current temperature for ""sensor \"%s\" on shelf \"%s\" at %s is %d C"
12338	Information	Temperature	"System temperature is OK" "System temperature is now ""within range: current temperature for sensor \"%s\" on shelf \"%s\" at %s is ""%d C"
12339	Warning	Temperature	"Temperature sensor failed" "Temperature sensor \"%s\" ""on shelf \"%s\" at %s has failed"
12340	Information	Temperature	"Temperature sensor is OK" "Temperature sensor \"%s\" ""on shelf \"%s\" at %s is operational: current temperature is %d C"
12341	Information	Cooling Fan	"Fan is at high speed" "Fan \"%s\" on shelf \"%s\" at %s %s is running ""fast: current speed %d rpm"
12342	Information	Cooling Fan	"Fan is at low speed" "Fan \"%s\" on shelf \"%s\" at %s %s is running ""slow: current speed %d rpm"
12343	Information	Cooling Fan	"Fan is OK" "Fan \"%s\" on shelf \"%s\" at %s %s is now running ""at speed %d rpm"
12344	Warning	Cooling Fan	"Fan stopped" "Fan \"%s\" on shelf \"%s\" at %s %s stopped working"
12345	Warning	Cooling Fan	"Fan missing" "Fan \"%s\" on shelf \"%s\" at %s %s is missing"
12346	Warning	Power Supply	"Power supply fail" "Power supply \"%s\" on shelf \"%s\" at %s is ""not connected or has failed"

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12347	Warning	Power Supply	"Power supply missing" "Power supply \"%s\" on shelf \"%s\" at %s is ""missing"
12348	Information	Power Supply	"Power supply is OK" "Power supply \"%s\" on shelf \"%s\" at %s is OK"
12601	Warning	NVRAM	"NVRAM battery is charging" "NVRAM battery is charging on controller %s"
12602	Information	NVRAM	"NVRAM battery has finished charging" "NVRAM battery has finished charging on controller %s"
12603	Critical	NVRAM	"NVRAM version is incompatible" "NVRAM version is incompatible with system firmware on controller %s"
12604	Error	NVRAM	"NVRAM card multi-bit error (MBE) detected" "The NVRAM memory content is corrupted on controller %s"
12605	Warning	NVRAM	"NVRAM card single-bit error (SBE) detected" "The NVRAM memory content is experiencing bit flips on controller %s"
12606	Error	NVRAM	"NVRAM card non-correctable multi-bit error (MBE) detected" "The NVRAM memory content has experienced a non-correctable bit flip error on controller %s"
12607	Warning	NVRAM	"NVRAM card correctable single-bit error (SBE) detected" "The NVRAM memory content has experienced a correctable bit flip error on controller %s"
12608	Information	NVRAM	"NVDIMM reserved flash blocks reached the upper threshold" "The reserved flash blocks have reached the upper threshold (remaining = %d) on controller %s."
12609	Warning	NVRAM	"NVDIMM reserved flash blocks reached the lower threshold" "The reserved flash blocks have reached the lower threshold (remaining = %d) on controller %s. Please contact Nimble Storage Support."

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12610	Warning	NVRAM	"NVDIMM ultracap cacapcitance discharged" "The NVDIMM ultracap capacitance is discharged (remaining = %d) on controller %s"
12611	Critical	NTB	"NTB_BAD_LINK" "Controller %s %s in %s is higher than the threshold value of %d. Please Contact Nimble Storage Support."
12612	Warning	NVRAM	"NVDIMM ultracap capacitance discharged" "The NVDIMM ultracap capacitance is discharged (remaining = %d) on controller %s. Please Contact Nimble Storage Support."
12613	Critical	NVRAM	"NVRAM module missing" "NVRAM module is not detected on controller %s. Please contact Nimble Storage Support."
12614	Critical	NVRAM	"NVDIMM ultracap capacitance discharged" "The NVDIMM ultracap capacitance is discharged on controller %s. Data is unsafe if power is lost. ""If this situation continues for %d minutes
12615	Critical	NVRAM	"Access to encrypted volumes denied" "Access to encrypted volumes denied. You must re-enter the passphrase to access the encrypted volumes. If you have swapped out both controllers
12701	Warning	Shelf	"Shelf controller connected to wrong side of host controller" "Controller %s of Shelf location % s
12702	Error	Shelf	"Shelf SES device not ready" "SES device not ready for shelf serial %s location %s"
12703	Warning	Shelf	"SAS link of a shelf slot is degraded" "SAS link of shelf serial %s location %s slot %s is degraded"
12704	Warning	Shelf	"Not enough disk in shelf" "Shelf serial %s location %s does not have enough number of disks"
12705	Error	Shelf	"System cannot read shelf serial number" "System cannot read the serial number from shelf at location %s"

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12706	Error	Shelf	"SAS cable connected to wrong port" "SAS cable s) that connected to shelf location %s is(are) connected to the wrong port(s
12707	Warning	Shelf	"SAS link(s) between the shelves is(are) degraded" "SAS link(s) between shelf location %s and %s is(are) degraded"
12708	Warning	Shelf	"SAS expander error" "SAS expander in shelf at location % s
12709	Error	Shelf	"Shelf SAS loop detected" "Both shelf controllers A and B are connected to the same head controller at location %s and %s. ""Ensure that each controller on the shelf is connected to its head unit equivalent (A or B)"
12710	Warning	Shelf	"Shelf appear only on one side of host controllers" "Shelf serial %s visible on host controller %s location %s but not on host controller % s
12711	Warning	Shelf	"Shelf controller connection order mismatched" "Shelf serial %s is at location %s on host controller A but at location %s on host controller B. ""Connect both controllers on each shelf to the head or shelf in the same order"
12712	Error	Shelf	"Cannot access shelf SES device" "Cannot access shelf SES device in the shelf serial %s at location % s
12713	Warning	Shelf	"Controller failover due to shelf problem" "Controller failover due to shelf problem"
12714	Information	Shelf	"New Shelf detected" "A new shelf is detected at location %s."
12715	Information	Shelf	"Shelf is disconnected " "Shelf at location % s
12716	Warning	Shelf	"Shelf chassis swap detected" "Shelf serial %s location %s replaced shelf serial %s"
12717	Warning	Shelf	"Shelf daisy chain too long" "Shelf serial %s location %s exceeds maximum supported daisy chain length"

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12719	Information	Shelf	"Connection to shelf is restored" "SAS connection to shelf at location % s
12720	Warning	Shelf	"SAS link disabled" "SAS phy %s on shelf at location %s has excessive link error and has been disabled"
12721	Warning	Shelf	"Excessive link error on disk" "SAS link between disk sn %s and expander %s has excessive link error"
12730	Error	Shelf	"SAS cable connected to wrong port" "SAS cable s) that connected to shelf location %s is(are) connected to the wrong port(s
12731	Error	Shelf	"Cannot access shelf SES device" "Cannot access shelf SES device in the shelf serial %s at location % s
12732	Warning	Shelf	"SAS link disabled" "SAS phy %s on shelf at location %s has excessive link errors and has been disabled."
12733	Warning	Shelf	"Excessive link error on disk" "SAS link between disk S/N %s and expander %s has excessive link errors. Contact Nimble Storage Support."
12734	Warning	Shelf	"SAS cable connected to wrong port" "SAS cable s) that connected to shelf location %s is(are) connected to the wrong port(s
12735	Warning	Shelf	"Cannot access shelf SES device" "Cannot access shelf SES device in the shelf serial %s at location % s
12736	Warning	Shelf	"Cannot access shelf SES device" "Cannot access shelf SES device in the shelf serial %s at location % s
12738	Warning	Shelf	"Cannot access shelf SES device" "Cannot access shelf SES device in the shelf with S/N %s at location %s. Some sensor data ""may not be available. Contact Nimble Storage Support."

Hardware

Alert/Event	Severity	Type	Description/Possible Cause/Actions
12740	Warning	Shelf	"Cannot access shelf IPMI device" "Cannot access shelf IPMI device in the shelf with S/N %s at location %s. Some sensor data ""may not be available. Contact Nimble Storage Support."
12741	Warning	Shelf	"SAS cable connected to wrong port" "SAS cable between shelf location %s and location %s is connected to wrong port."
12742	Warning	Shelf	"Excessive link errors between disk and expander" "There are excessive link errors on the SAS link between disk slot %s and expander ""in the shelf %s. Contact Nimble Storage Support."
12743	Warning	Shelf	"Shelf bad interconnect phy" "There are excessive link errors on the SAS link between %s and %s. ""Contact Nimble Storage Support."
12744	Warning	Shelf	"Excessive link errors between shelves" "There are excessive link errors on the SAS link between %s and %s. ""Contact Nimble Storage Support."
12901	Critical	Controller	"Physical memory detected is less than installed" "Physical memory detected is less than installed on array serial number %s controller %s. Installed %ld GB
14000	Information	Array	"CS-Model changed" "The CS-Model has changed. The model is now %s."
14001	Warning	Array	"CS-Model unknown" "The system cannot determine the CS-Model. The prior model was % s
14002	Warning	Temperature	"The system temperature is beginning to exceed the allowable operating ""temperature. If it continues to exceed allowed operating temperature ""the system may shut down."
14003	Critical	Temperature	"Controller shutdown occurred due to excessive temperature" "Temperature exceeds maximum operating temperature (50 Celsius) on controller % s

Hardware

Alert/Event	Severity	Type	Description/Possible Cause/Actions
14004	Critical	Temperature	"Controller exceeds allowable operating temperature" "Temperature exceeds allowable operating temperature on controller %s (%d Celsius)"
14005	Critical	Temperature	"Backplane exceeds allowable operating temperature" "Temperature exceeds allowable operating temperature on backplane (%d Celsius)"
14006	Warning	Temperature	"System temperature is too high." "Sensor %s on shelf %s at %s is %d Celsius. ""If any sensor exceeds 50 degrees C
14007	Critical	Temperature	"System temperature continues to be high." "Temperature Sensor on shelf %s is %d Celsius. ""If any sensor exceeds 50 degrees C
14008	Critical	Power Supply	"Power supply down revisioned" "While your shelf %s continues to function normally
14009	Critical	Temperature	"Controller shutdown occurred due to excessive temperature" "Temperature exceeds maximum operating temperature (%lu) on controller % s
14010	Critical	Power Supply	"Power supply down revisioned" "While your shelf %s continues to function normally
14011	Warning	Temperature	"Array temperature is too high" "Temperature sensor %s on shelf %s at %s is %d Celsius. ""Check air temperature and air flow around the array."
14012	Critical	Temperature	"Array temperature continues to be high" "Temperature sensor on shelf %s is %d Celsius. ""The shelf shuts down if it exceeds the maximum operating temperature."
14013	Critical	Power Supply	"Power supply status unavailable" "Array software is unable to query status of power supply units (PSUs). Although the shelf S/N %s continues to function normally

Hardware			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
14014	Critical	Power Supply	"Power supply status unavailable" "Unable to query status of power supply units (PSUs). Although the shelf S/N %s continues to function normally
14015	Critical	Temperature	"Array temperature continues to be high" "Temperature sensor on shelf %s is %d Celsius which exceeds the safe operating temperature of %d Celsius. ""Do not operate shelf under elevated temperatures. To ensure data protection
14200	Error	Controller	"An unknown firmware version was detected" "An unknown firmware version was detected for component %s on controller %s"
14400	Information	FC	"Fibre Channel link up" "Fibre Channel link up on controller %s port %s"
14401	Warning	FC	"Fibre Channel link down" "Fibre Channel link down on controller %s port %s"
14402	Information	FC	"Fibre Channel link not connected" "Fibre Channel link not connected on controller %s port %s"
14403	Information	FC	"Fibre Channel link up with connection to fabric" "Fibre Channel link up with connection to fabric on controller %s port %s."
14404	Warning	FC	"Fibre Channel link up with no connection to fabric" "Fibre Channel link up with no connection to fabric on controller %s port %s. Check the switch log files for errors."

Table 16: Alert Replication

Replication			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
11000	Information	Partner	"Replication succeeded" "Successfully replicated snapshot collection %s ""to partner %s"
11001	Information	Partner	"Excessive replication delay" "Excessive delay replicating volume collection %s ""to partner %s"

Replication			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
11002	Warning	Partner	"Excessive replication delay" "Excessive delay replicating volume collection %s ""to partner %s"
11003	Warning	Partner	"Partner synchronization failed" "Failed to synchronize replication configuration ""to partner % s
11004	Warning	Partner	"Volume replication stalled" "Replication is stalled. Volume %s on partner %s must be removed ""in order to proceed"
11005	Information	Partner	"Volume collection handover completed" "Completed handover of volume collection %s"
11006	Warning	Partner	"Replication between multi-array scaleout group"" and pre-scaleout array is not supported" "Cannot replicate from/to the pre-2.x partner %s from a multi-array group"
11007	Warning	Partner	"Replication between a scaleout group"" and pre-scaleout array must be configured over the management network" "Cannot replicate from/to the pre-2.x partner %s over the data network."
11008	Error	Partner	"Replication resynchronized" "Successfully resynchronized while replicating snapshot %s on volume %s ""from partner %s. Prior snapshots for this volume may be out of sync. ""Please contact Nimble Storage Support."
11009	Warning	Partner	"Resynchronization not supported" "Downstream partner %s has requested resynchronization of replicated snapshots"" for volume %s. Please contact Nimble Storage Support for assistance."
11010	Warning	Partner	"Resynchronization not supported" "Volume %s needs resynchronization from upstream partner %s. ""Contact Nimble Storage Support."
11011	Error	Partner	"Replication resynchronized" "Successfully resynchronized while replicating snapshot %s on volume %s ""from partner %s. Prior snapshots for this volume may be out of sync. ""Contact Nimble Storage Support."

Replication			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
11012	Warning	Partner	"Resynchronization not supported" "Downstream partner %s has requested resynchronization of replicated snapshots"" for volume %s. Contact Nimble Storage Support."
11013	Warning	Partner	"Partner authentication failed" "Failed to authenticate with replication partner %s because partner secrets do not match."" Replication is currently paused. Edit the partner's shared secret to resume replication."
11014	Warning	Partner	"Replicated snapshot resynchronized" "Successfully resynchronized replicated snapshot %s of volume %s ""from partner %s. Prior snapshots for this volume may not be synchronized. ""Contact Nimble Storage Support."
11015	Warning	Partner	"Replicated snapshot resynchronization not supported" "Downstream partner %s has requested resynchronization of replicated snapshots"" for volume %s. This group does not support resynchronization. Contact Nimble Storage Support."
11016	Warning	Partner	"Replication partner authentication failed" "Failed to authenticate with replication partner %s because shared secret does not match."" Replication is currently paused. Edit partner and match the shared secret to resume replication."
11017	Error	Volume	"Replicated volume needs resynchronization" "Snapshot %s of volume %s needs resynchronization from upstream partner. ""Contact Nimble Storage Support."
11018	Information	Partner	"Volume collection handover aborted" "Aborted handover of volume collection %s"
11019	Information	Partner	"Volume collection handover completed with errors." "Completed handover of volume collection % s
11020	Information	Partner	"Volume collection handover aborted with errors." "Aborted handover of volume collection % s

Replication			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
11021	Information	Partner	"Volume collection could not be deleted on the downstream replication partner." "Volume collection %s could not be deleted on the downstream replication partner and must be manually removed."
11022	Warning	Volume	"Replication of volume is stalled due to space reclamation." "Replication is stalled for volume %s of volume collection %s. Volume must be deleted on the downstream for replication to continue."
11023	Warning	Volume	"Replication of volume is stalled due to space reclamation on common snapshot." "Replication is stalled for volume %s of volume collection % s
11024	Warning	Volume	"Replicated volume needs resynchronization" "Replicated volume %s needs resynchronization from upstream partner. ""Resynchronization on volume %s has been initiated."
11025	Information	Volume	"Replicated volume resynchronized" "Successfully resynchronized replicated volume %s from partner %s."

Table 17: Alert Security

Security			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
14806	Warning	Controller	"Root login succeeded." "Root login to controller %s from %s succeeded."
14807	Critical	Controller	"Root login failed." "Root login to controller %s from %s failed."
14808	Warning	Controller	"Login by nsupport succeeded" "Login by nsupport to controller %s from %s succeeded."
14809	Critical	Controller	"Login by nsupport failed" "Login by nsupport to controller %s from %s failed."
14810	Warning	Controller	"Elevating user privilege to root succeeded" "Elevating user privilege from %s to root on controller %s succeeded."

Security			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
14811	Critical	Controller	"Elevating user privilege to root failed" "Elevating user privilege from %s to root on controller %s failed."
14812	Note	Group	"Login attempt by user without an assigned role" "Login attempt by Active Directory user %s failed. The Active Directory group to which user belongs to is not associated with a role on the Nimble Storage group %s. Associate an Active Directory group the user belongs to with a role on Nimble Storage group. Ensure that the role has appropriate privileges before attempting to log in again."

Table 18: Alert Service

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
0101	Warning	Array	"Low writable space" "The system is low on writable space. Write operations will be slowed. "
0102	Information	Array	"Writable space reclaimed" "Writable space is sufficient
0103	Critical	Array	"Critically low writable space" "The system is extremely low on writable space
0104	Warning	Array	"High system utilization" "The system is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
0105	Information	Array	"System utilization OK" "The system is down to %d%% utilization."
0106	Critical	Array	"Critically high system utilization" "The system is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
0107	Critical	Array	"Critically low writable space" "The system's writable space is extremely low. Application writes will be progressively slowed as writable space continues to decline."

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
0110	Critical	Array	"Critically low writable space" "The system's writable space is too low for write operations. Application writes will time out."
0111	Critical	Array	"Critically low writable space" "The array's writable space is extremely low. Application writes will be progressively slowed as writable space continues to decline."
0112	Critical	Array	"Critically low writable space" "The array's writable space is too low for write operations. Application writes will time out."
0113	Critical	Array	"Critically high system utilization" "The system is extremely low on space. Incoming writes will be denied."
0114	Critical	Array	"Critically low space" "The space on the array is extremely low. Application writes will be progressively slowed as writable space continues to decline."
0115	Information	Array	"Writable space sufficiently restored" "The space on the array has been sufficiently restored; write operations are back to normal speed."
0116	Critical	Array	"Critically low space" "The space on the array is too low for write operations. Application writes will time out."
2001	Critical	Controller	"Unhandled controller exception" "Unhandled exception on controller %s"
2002	Information	Controller	"Controller takeover occurred" "Takeover occurred on controller %s"
2003	Information	Controller	"Standby controller available" "Standby controller %s available"
2004	Warning	Controller	"Standby controller not available" "Standby controller %s not available"
2005	Critical	Controller	"Excessive controller restarts detected" "Excessive controller %s restarts detected"
2006	Critical	Controller	"Restarting controller to recover service" "Restarting controller %s to recover %s service"

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
2007	Information	Controller	"Rebooting controller on user request" "Rebooting controller %s on user request"
2008	Error	Controller	"Failed to reboot controller per user request" "Failed to reboot controller %s per user request"
2009	Information	Controller	"Halting controller on user request" "Halting controller %s on user request"
2010	Error	Controller	"Failed to halt controller per user request" "Failed to halt controller %s per user request"
2011	Information	Controller	"Standby controller not available" "Standby controller %s not available"
2012	Warning	Controller	"Standby controller not available for an extended period" "Standby controller %s not available for an extended period"
2013	Warning	Controller	"Unhandled controller exception" "Unhandled exception on controller %s"
2014	Critical	Controller	"Unhandled controller exception all services are down"
2015	Warning	Controller	"Standby controller not available for an extended period" "Standby controller %s not available for %d minutes"
2016	Critical	Controller	"Standby controller not available for an extended period" "Standby controller %s not available for %d minutes"
2017	Warning	Controller	"Controller takeover occurred" "Takeover occurred on controller %s"
2018	Information	Controller	"Controller failover occurred" "Failover occurred on controller %s"
2019	Information	Controller	"Controller failover occurred" "Failover of active role occurred on controller %s"

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
2020	Critical	Controller	"Controller running with watchdog disabled" "Watchdog disabled on controller %s. Disconnect all connected clients and contact Nimble Storage Support immediately."
2021	Information	Controller	"Controller failover occurred" "Failover of active role occurred on controller %s. Controller %s is now the active controller."
2025	Critical	Controller	"Standby controller missing too many drives and cannot failover" "Shelf serial %s standby controller %s missing too many drives and cannot failover."
2026	Warning	Controller	"Standby controller performance degraded on failover; missing afs" "Shelf serial %s standby controller %s performance degraded on failover; missing afs."
2027	Critical	Controller	"Excessive controller restarts detected" "Excessive controller restarts detected. Controller %s is starting up in degraded mode. Some processes will be disabled. Contact Nimble Storage Support."
2101	Information	Service	"Service started" "System services started. Software version is %s"
2102	Information	Service	"Service stopped unexpectedly restarting"
2103	Warning	Service	"Service stopped unexpectedly not restarting"
2104	Information	Service	"Created secure tunnel to Nimble Storage support" "Created secure tunnel to Nimble Storage support on controller %s"
2105	Information	Service	"Tunnel to Nimble Storage support has been terminated" "Tunnel to Nimble Storage support has been terminated on controller %s"
2106	Warning	Service	"Service stopped unexpectedly" "The %s stopped unexpectedly on the array

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
2107	Warning	Service	"Service stopped unexpectedly" "The %s stopped unexpectedly on the array
2108	Warning	Service	"Failed to send alert e-mail" "Failed to send e-mail to % s
2109	Warning	Service	"Failed to send alert e-mail" "Failed to send alert e-mail to SMTP server. Verify SMTP server and port configuration."
2110	Warning	Service	"Service manually disabled" "The %s was manually disabled on an array controller and will not be started."
2111	Information	Service	"Created secure tunnel to Nimble Storage Support" "Created secure tunnel to Nimble Storage Support on controller %s."
2112	Information	Service	"Tunnel to Nimble Storage Support has been terminated" "Tunnel to Nimble Storage Support has been terminated on controller %s."
10010	Critical	Array	"A member array became unreachable" "Member array %s became unreachable"
10011	Information	Array	"A member array became reachable" "Member array %s became reachable"
10012	Critical	Array	"A member array is unreachable" "Member array %s is unreachable"
10013	Information	Array	"A member array is reachable" "Member array %s is reachable"
10229	Error	Pool	"Failed to assign array to pool" "Failed to assign array %s to pool. Please retry after verifying that array is up and has network connectivity."
10230	Error	Pool	"Failed to unassign array from pool" "Failed to unassign array %s from pool. Please retry after verifying that array is up and has network connectivity."

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10240	Error	Pool	"Failed to merge pools." "Failed to merge pool %s into pool. Please retry after verifying that all arrays in both pools are up and have network connectivity."
10252	Information	Pool	"Completed unassign of array from pool" "Completed unassign of array %s from pool. All data on the array has been moved to other arrays in the pool. The array can now be assigned to other pools or removed from the group."
10253	Information	Pool	"Completed unassign of array from pool" "Completed unassign of array %s from pool %s. All data on the array has been moved to other arrays in the pool. The array can now be assigned to other pools or removed from the group."
10256	Warning	Pool	"Failed to assign array to pool" "Failed to assign array %s to pool %s. Retry after verifying that array is up and has network connectivity."
10257	Warning	Pool	"Failed to unassign array from pool" "Failed to unassign array %s from pool %s. Retry after verifying that array is up and has network connectivity."
10260	Error	Pool	"Failed to merge pools." "Failed to merge pool %s into pool %s. Retry after verifying that all arrays in both pools are up and have network connectivity."
10282	Critical	Volume	"Insufficient cache capacity in the pool for cache pinned volumes" "Pinnable cache capacity %d MB in pool %s is insufficient to keep all cache pinned volumes in cache. Consider adding more cache or unpinning some volumes to restore performance guarantees for cache pinned volumes."
10401	Warning	Array	"Used space above warning limit" "The amount of used space has reached %d%% of free space"
10402	Information	Array	"Used space below warning limits" "The amount of used space is below %d%% of free space"

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10403	Critical	Array	"Used space above critical level" "The amount of used space has reached a critical level at %d%% of free space. Volumes above their reserve will be taken offline once the free space has been exhausted"
10404	Information	Array	"Used space below critical limits" "The amount of used space is at %d%% of free space. Volumes above their reserve will be taken offline once the free space has been exhausted"
10405	Information	Array	"Array space utilization OK" "The Array is down to %d%% space utilization."
10406	Information	Array	"High array space utilization" "The array space is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
10407	Warning	Array	"High array space utilization" "The array space is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
10408	Critical	Array	"Critically high array space utilization" "The array space is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
10409	Warning	Pool	"Used space in the pool above warning limit" "The amount of used space in the pool has reached %d%% of free space"
10410	Information	Pool	"Used space in the pool below warning limits" "The amount of used space in the pool is below %d%% of free space"
10411	Critical	Pool	"Used space in the pool above critical level" "The amount of used space in the pool has reached a critical level at %d%% of free space. Volumes above their reserve will be taken offline once the free space has been exhausted"
10412	Information	Pool	"Used space in the pool below critical limits" "The amount of used space in the pool is at %d%% of free space. Volumes above their reserve will be taken offline once the free space has been exhausted"

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10413	Information	Pool	"Pool space utilization OK" "The pool is down to %d%% space utilization."
10414	Information	Pool	"High pool space utilization" "The pool space is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
10415	Warning	Pool	"High pool space utilization" "The pool space is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
10416	Critical	Pool	"Critically high pool space utilization" "The pool space is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
10417	Warning	Pool	"Used space in the pool above warning limit" "The amount of used space in the pool %s has reached %d%% of free space."
10418	Information	Pool	"Used space in the pool below warning limits" "The amount of used space in the pool %s is below %d%% of free space."
10419	Critical	Pool	"Used space in the pool above critical level" "The amount of used space in the pool %s has reached a critical level at %d%% of free space. Volumes above their reserve will be taken offline once the free space has been exhausted."
10420	Information	Pool	"Used space in the pool below critical limits" "The amount of used space in the pool %s is at %d%% of free space. Volumes above their reserve will be taken offline once the free space has been exhausted."
10421	Information	Pool	"Pool space utilization OK" "The pool %s is down to %d%% space utilization."
10422	Information	Pool	"High pool space utilization" "The pool space in pool %s is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10423	Warning	Pool	"High pool space utilization" "The pool space in pool %s is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
10424	Critical	Pool	"Critically high pool space utilization" "The pool space in pool %s is currently %d%% full. Consider deleting unneeded snapshots or volumes to reduce space usage."
12400	Warning	Raid	"RAID array degraded" "RAID array degraded by %d %s"
12401	Information	Raid	"RAID array rebuild started" "RAID array started rebuilding disk %s"
12402	Information	Raid	"RAID array rebuild is done" "RAID array has successfully ""completed rebuilding disk %s"
12403	Warning	Raid	"RAID array rebuild failed" "Failed to rebuild RAID array on ""disk % s
12404	Warning	Raid	"RAID array rebuild failed" "Failed to rebuild RAID array on ""disk % s
12405	Warning	Raid	"Disks missing from RAID array" "%d %s missing from RAID ""array"
12406	Information	Raid	"Disk marked as spare" "Disk %s is marked as spare"
12407	Critical	Raid	"Could not assemble RAID array" "Could not assemble RAID array
12408	Warning	Raid	"RAID array degraded" "RAID array degraded by %d %s on %s shelf %s"
12409	Information	Raid	"RAID array rebuild started" "RAID array started rebuilding disk %s ""on %s shelf %s"
12410	Information	Raid	"RAID array rebuild is done" "RAID array has successfully ""completed rebuilding disk %s on %s shelf %s"

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12411	Warning	Raid	"RAID array rebuild failed" "Failed to rebuild RAID array from ""%s shelf %s on disk % s
12412	Warning	Raid	"RAID array rebuild failed" "Failed to rebuild RAID array from ""%s shelf %s on disk % s
12413	Warning	Raid	"Disks missing from RAID array" "%d %s missing from RAID ""array on %s shelf %s"
12414	Information	Raid	"Disk marked as spare" "Disk %s is marked as spare on %s shelf %s"
12415	Critical	Raid	"Could not assemble RAID array" "Could not assemble RAID array on ""%s shelf % s
12416	Information	Raid	"RAID array rebuild scheduled" "RAID array rebuild scheduled on disk %s ""of %s shelf %s"
12417	Information	Raid	"RAID array rebuild stopped" "RAID array rebuild stopped on disk %s ""of %s shelf %s"
12418	Information	Raid	"RAID array redundancy" "RAID array on %s shelf %s now has %d/%d disks
12501	Warning	iSCSI	"Too many iscsi connections" "iscsi login rejected from %s"
12502	Warning	iSCSI	"Too many unaligned iscsi reads/writes" "unaligned reads/writes percentage %.2f%% "
12504	Warning	iSCSI	"Too many iSCSI connections" "iSCSI login rejected from %s"
12505	Warning	iSCSI	"Too many unaligned iSCSI reads/writes" "unaligned reads/writes percentage %.2f%% "
12506	Warning	iSCSI	"Too many iSCSI connections" "iSCSI login from %s is rejected"
12718	Information	Shelf	"New shelf is activated" "Shelf serial %s at location %d has been activated"

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
12737	Information	Shelf	"New shelf is activated" "Shelf serial %s at location %s has been activated"
12739	Information	Shelf	"New shelf activated" "Shelf with S/N %s at location %s was activated"
13501	Information	Array	"Group Leader role takeover succeeded" "Successfully takenover group leader role from array '%s' (serial: %s)"
13504	Information	Array	"Group Leader role takeover succeeded" "Array %s (serial: %s) successfully took over group leader role from array %s (serial: %s)"
13505	Information	Array	"Group Leader role take over succeeded" "Array %s S/N %s is now the group leader. Previous group leader was array %s S/N %s."
13506	Error	Array	"Group Leader role migrate failed" "Array %s S/N: %s failed to migrate group leader role to array %s S/N: %s."
13502	Error	Array	"Group Leader role takeover has been rejected by a member array" "Group Leader role takeover has been rejectec by member array '%s'"
13503	Error	Array	"Group Leader role takeover has been rejected by a member array" "Group Leader role takeover has been rejected by member array '%s'"
13601	Information	Group	"Group merge completed successfully" "Group merge with group %s completed successfully"
13602	Warning	Group	"Failed to stop services on source group for group merge" "Failed to stop services on source group %s for group merge. Group merge operation would be rolled back. Retry the operation later."
13603	Information	Group	"Completed rollback of group merge operation" "Completed rollback of group merge operation with group %s. Services on group %s have been resumed."

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
13604	Warning	Group	"Group merge operation is not complete on some arrays operation would be retried"
13605	Warning	Group	"Failed to merge configuration from source group" "Failed to merge configuration from source group %s. Group merge operation would be rolled back. Retry the operation later."
13701	Warning	Group	"Data Migration is getting delayed because of continuous restarts" "Data Migration is getting delayed because of continuous restarts. Contact Nimble Storage Support."
13702	Warning	Group	"Data migration is delayed because of repeated restarts" "Data migration is delayed because of repeated restarts. Contact Nimble Storage Support."
13703	Warning	Group	"Data migration is delayed because of repeated errors" "Data migration is delayed because of repeated errors. Contact Nimble Storage Support."
14700	Warning	Service	"Oldest events removed from the system" "Total number of events reached %lu. Maximum retention count is %lu. ""The oldest %lu events removed."
14701	Warning	Service	"Events count reached warning threshold" "Events count reached warning threshold of %lu events. If the number of events ""reaches the maximum retention count of % lu
14702	Warning	Service	"Oldest audit log records removed from the system" "Total number of audit log records reached %lu. Maximum retention count is %lu. ""The oldest %lu audit log records removed."
14703	Warning	Service	"Audit log records count reached warning threshold" "Audit log records count reached warning threshold of %lu. If the number of audit log records ""reaches the maximum retention count of % lu

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
14704	Information	Service	"Oldest events removed from the system" "Total number of events reached %lu. Maximum retention count is %lu. ""The oldest %lu events removed."
14705	Information	Service	"Events count reached warning threshold" "Events count reached warning threshold of %lu events. If the number of events ""reaches the maximum retention count of % lu
14706	Information	Service	"Oldest audit log records removed from the system" "Total number of audit log records reached %lu. Maximum retention count is %lu. ""The oldest %lu audit log records removed."
14707	Information	Service	"Audit log records count reached warning threshold" "Audit log records count reached warning threshold of %lu. If the number of audit log records ""reaches the maximum retention count of % lu
14800	Warning	Controller	"Root Login succeeded." "Root Login to controller %s from %s succeeded."
14801	Critical	Controller	"Root Login failed." "Root Login to controller %s from %s failed."
14802	Warning	Group	"Active Directory client service on this group is not running" "Active Directory client service on this group is not running. To restart the service
14803	Warning	Group	"Active Directory Domain Controller is not reachable" "Active Directory Domain Controller for domain %s is not reachable. Contact your Active Directory administrator."
14804	Warning	Group	"Could not authenticate with Active Directory Domain Controller" "Could not authenticate with Active Directory Domain Controller due to invalid credentials. Remove this group from the domain and retry join."

Service			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
14805	Information	Group	"Successfully communicated with Active Directory Domain Controller" "Successfully communicated with Active Directory Domain Controller."

Table 19: Alert Test

Test			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
5000	DBG	Test	"Test alert" "Test message at debug level"
5001	Information	Test	"Test alert" "Test message at info level"
5002	Error	Test	"Test alert" "Test message at error level"
5003	Note	Test	"Test alert" "Test message at \"not\" level"
5004	Note	Test	"Test alert" "Test message at note level"
5005	Warning	Test	"Test alert" "Test message at warn level"
5006	Critical	Test	"Test alert" "Test message at critical level"
5007	Note	Test	"Test alert" "Test message at notice level"
5008	Information	Test	"Test alert" "Test recovery from warning alarm message"
5009	Information	Test	"Test alert" "Test recovery from critical alarm message"

Table 20: Alert Unknown

Unknown			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10101	Information	Array	"Cache cleared by user" "Cache cleared by user"

Table 21: Alert Update

Update			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
6000	Error	Controller	"Unpacking update package failed" "Unpacking update package failed on controller %s: %s"
6001	Information	Controller	"Software update started" "Software update to version %s started on controller %s"
6002	Error	Controller	"Reverting software to previous version" "Reverting software to previous version %s on controller %s"
6003	Information	Controller	"Successfully updated software" "Successfully updated software to version %s on controller %s"
6004	Information	Controller	"Rolling back software" "Rolling back software to version %s on controller %s"
6005	Error	Controller	"Update package pre-check failed" "Update package Pre-check failed on controller %s"
6007	Error	Array	"Software update failed" "Software update failed: %s"
6008	Information	Controller	"Unpacking update package" "Unpacking version %s update package on controller %s"
6009	Information	Controller	"Unpacked update package" "Unpacked version %s update package on controller %s"
6010	Information	Controller	"Software update reboot" "%s on controller %s"

Update			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
6011	Error	Array	"Failed to download software update package" "Failed to download software update package: %s"
6012	Error	Controller	"Software update failed. \"/tmp\" is full" "Software update failed. \"/tmp\" on controller %s is full"
6013	Error	Controller	"Software update failed. scratch space is full" "Software update failed. scratch space on controller %s is full"
6014	Error	Controller	"Software update failed. \"/var\" is full" "Software update failed. \"/var\" on controller %s is full"
6015	Error	Controller	"Software update failed. Configuration space is full" "Software update failed. Configuration space on controller %s is full"
6016	Error	Controller	"Software update failed. Recovery OS space is full" "Software update failed. Recovery OS space is full on controller %s is full"
6017	Error	Controller	"Software update package was not found" "Software update package %s was not found on controller %s"
6018	Error	Controller	"Software update package has wrong signature" "Software update package %s on %s controller %s has wrong signature"
6019	Error	Controller	"Software update package has wrong checksum" "Software update package %s on %s controller has wrong checksum"
6020	Error	Controller	"Software update precheck failed. Network connectivity will degrade after software update" "Software udpate precheck failed. Network connectivity will degrade on controller %s after software update"

Update			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
6021	Error	Controller	"Software update failed. Update cannot be applied while services cannot be failed over to other controller" "Software update failed. Update cannot be applied while services cannot be failed over to other controller. ""Controller %s is not in standby mode and cannot take over services."
6022	Error	Controller	"Software update cannot be applied." "Software update cannot be applied. Controller %s is not in standby mode and cannot take over services that are required for a successful update."
6023	Error	Controller	"Software update precheck failed. Network connectivity will degrade after software update" "Software update precheck failed. Network connectivity will degrade on controller %s after software update"
6024	Error	Controller	"Software update precheck failed. RAID assembly failed on standby controller and prevents software update to proceed" "Software update precheck failed. RAID assembly failed on standby controller %s and prevents software update to proceed"
6025	Note	Controller	"Software update started" "Software update to version %s started on controller %s"
6026	Note	Controller	"Successfully updated software" "Successfully updated software to version %s on controller %s"
6401	Warning	Array	"Software download failed due to DNS errors" "Software download failed due to DNS lookup failure on controller %s for host %s"
6501	Warning	Array	"Failed to contact a member array during software update" "Failed to contact member %s during software update"
6502	Warning	Array	"Failed to update software on a member array" "Failed to update software on member array %s"
6503	Warning	Array	"Software update on a member array timed out" "Software update on member array %s timed out after %d minutes"

Update			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
6504	Note	Array	"Removed conflicting IQN that was added during software upgrade" "Removed conflicting IQN %s that was added during software upgrade."

Table 22: Alert Volume

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10200	Warning	Volume	"Volume space usage over warning limit" "Volume %s space usage is over the configured warning limit"
10201	Information	Volume	"Volume space usage below warning limit" "Volume %s space usage is now below the configured warning limit"
10202	Error	Volume	"Volume space usage over quota" "Volume %s space usage is over the configured quota"
10203	Information	Volume	"Volume space usage below quota" "Volume %s space usage is now below the configured quota"
10204	Warning	Volume	"Volume snapshot space usage over warning limit" "Volume %s snapshot space usage is over the configured warning limit"
10205	Information	Volume	"Volume snapshot space usage below warning limit" "Volume %s snapshot space usage is now below the configured warning limit"
10206	Error	Volume	"Volume snapshot space usage over quota" "Volume %s snapshot space usage is over the configured quota"
10207	Information	Volume	"Volume snapshot space usage below quota" "Volume %s snapshot space usage is now below the configured quota"
10208	Critical	Volume	"Volume space usage approaching quota" "Volume %s space usage is at %d%% and approaching quota of %d%%. It will be taken offline if it exceeds the quota"

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10209	Information	Volume	"Volume space usage below critical limit" "Volume %s space usage at %d%% is below critical limit
10210	Critical	Volume	"Volume snapshot space usage approaching quota" "Volume %s snapshot space usage is at %d%% and approaching quota. It will be taken offline if it exceeds the quota"
10211	Information	Volume	"Volume snapshot space usage below critical limit" "Volume %s snapshot space usage at %d%% is below critical limit
10212	Critical	Volume	"Volume being taken offline because it is above its quota" "Volume %s is being taken offline because it is above its quota"
10213	Warning	Volume	"The quota for volume is being set to 100%% because it is above its current quota" "The quota for volume %s is being set to 100%% because it is above its quota. This is a one time action. Adjust volume quotas appropriately as the system now enforces volume quotas by taking offline volumes with space usage above quota"
10214	Critical	Volume	"Volume being taken offline because it is above its snapshot quota" "Volume %s is being taken offline because it is above its snapshot quota"
10215	Warning	Volume	"The snapshot quota for volume is being set to unlimited because it is above its snapshot quota" "The snapshot quota for volume %s is being set to unlimited because it is above its snapshot quota. This is a one time action. Adjust snapshot quotas appropriately as the system now enforces volume snapshot quotas by taking offline volumes with snapshot space usage above quota"
10216	Critical	Volume	"Volume being taken offline because it is above its reserve and system is out of free space" "Volume %s is being taken offline because it is above its reserve and system is out of free space"

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10217	Warning	Volume	"Volume reserve is being set to zero because it is above its reserve and system is out of free space" "Volume %s reserve is being set to zero because it is above its reserve and system is out of free space. This is a one time action. Adjust volume reserves appropriately as the system now enforces volume reserves by taking offline volumes with space usage above their reserve when free space on the system is exhausted"
10218	Critical	Volume	"Volume being taken offline because it is above its snapshot reserve and system is out of free space" "Volume %s is being taken offline because it is above its snapshot reserve and system is out of free space"
10219	Warning	Volume	"Volume snapshot reserve is being set to zero because it is above its reserve and system is out of free space" "Volume %s snapshot reserve is being set to zero because it is above its reserve and system is out of free space. This is a one time action. Adjust volume snapshot reserves appropriately as the system now enforces snapshot reserves by taking offline volumes with snapshot space usage above reserve when system is out of free space"
10220	Warning	Volume	"Volume space usage over warning limit" "Volume %s space usage at %d%% is over the configured warning limit of %d%%"
10221	Information	Volume	"Volume space usage below warning limit" "Volume %s space usage at %d%% is now below the configured warning limit of %d%%"
10222	Critical	Volume	"Volume space usage is over the configured quota" "Volume %s space usage at %d%% is over the configured quota at %d%%"
10223	Information	Volume	"Volume space usage is now below the configured quota" "Volume %s space usage at %d%% is now below the configured warning limit of %d%%"

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10224	Warning	Volume	"Volume snapshot space usage over warning limit" "Volume %s snapshot space usage is over the configured warning limit"
10225	Information	Volume	"Volume snapshot space usage below warning limit" "Volume %s snapshot space usage is now below the configured warning limit"
10226	Critical	Volume	"Volume snapshot space usage over quota" "Volume %s snapshot space usage is over the configured quota"
10227	Information	Volume	"Volume snapshot space usage below quota" "Volume %s snapshot space usage is now below the configured quota"
10228	Warning	Volume	"Volume attributes synchronization to arrays in the pool delayed" "Volume %s attributes synchronization to arrays in the pool delayed. Will retry"
10231	Warning	Volume	"Volume space usage approaching quota" "Volume %s space usage is at %d%% and approaching quota of %d%%."
10232	Warning	Volume	"Volume snapshot space usage approaching quota" "Volume %s snapshot space usage is at %d%% and approaching quota."
10233	Warning	Volume	"Volume space usage is above its reserve" "Volume %s space usage at %d%% is above its reserve at %d%%. It will be taken offline once the free space has been exhausted."
10234	Information	Volume	"Volume space usage is below its reserve" "Volume %s space usage at %d%% is below its reserve at %d%%."
10235	Warning	Volume	"Volume snapshot space usage is above its reserve" "Volume %s snapshot space usage at %d%% is above its snapshot reserve at %d%%. It will be taken offline once the free space has been exhausted."

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10236	Information	Volume	"Volume snapshot space usage is below its reserve" "Volume %s snapshot space usage at %d%% is below its snapshot reserve at %d%%."
10241	Information	Volume	"Volume attributes synchronization to arrays in the pool succeeded" "Volume %s attributes synchronization to arrays in the pool succeeded."
10242	Warning	Volume	"Snapshot attributes synchronization to arrays in the pool delayed" "Snapshot %s in Volume %s attributes synchronization to arrays in the pool delayed. Will retry"
10243	Information	Volume	"Snapshot attributes synchronization to arrays in the pool succeeded" "Snapshot %s in Volume %s attributes synchronization to arrays in the pool succeeded."
10250	Information	Volume	"Volume completed move to destination pool" "Volume %s completed move to destination pool %s. Volume data is now on destination pool."
10251	Information	Volume	"Volume completed aborting move to destination pool" "Volume %s completed aborting move to destination pool %s. Volume data has been moved back to source pool %s."
10261	Critical	Volume	"Volume space usage approaching volume quota" "Volume %s space usage is at %d%% and approaching volume quota of %d%%. It will be set to non-writable if it exceeds the volume quota."
10262	Critical	Volume	"Volume being set to non-writable because it is above its volume quota" "Volume %s is being set to non-writable because it is above its volume quota."
10263	Critical	Volume	"Volume snapshot space usage approaching snapshot quota" "Volume %s snapshot space usage is at %d%% and approaching snapshot quota. It will be set to non-writable if it exceeds the snapshot quota."

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10264	Critical	Volume	"Volume being set to non-writable because it is above its snapshot quota" "Volume %s is being set to non-writable because it is above its snapshot quota."
10265	Critical	Volume	"Volume being set to non-writable because it is above its reserve and system is out of free space" "Volume %s is being set to non-writable because it is above its reserve and system is out of free space"
10266	Critical	Volume	"Volume being set to non-writable because it is above its snapshot reserve and system is out of free space" "Volume %s is being set to non-writable because it is above its snapshot reserve and system is out of free space"
10278	Critical	Volume	"Invalid volume serial number detected" "Volume %s is being taken offline because it is ""using an invalid serial number %s. Contact ""Nimble Storage Support."
10279	Critical	Volume	"Invalid snapshot serial number detected" "Snapshot %s of volume %s is being taken offline because it is ""using an invalid serial number %s. Contact Nimble Storage Support."
10283	Warning	Folder	"Folder attributes synchronization to arrays in the pool delayed" "Folder %s attributes synchronization to arrays in the pool delayed. Will retry"
10284	Information	Folder	"Folder attributes synchronization to arrays in the pool succeeded" "Folder %s attributes synchronization to arrays in the pool succeeded."
10300	Information	Protection Set	"Scheduled snapshot succeeded" "Successfully snapshoted volumes associated with ""volume collection %s schedule %s"
10301	Warning	Protection Set	"Scheduled snapshot failed" "Failed to snapshot volumes associated with volume collection %s ""schedule % s
10302	Information	Protection Set	"Scheduled snapshot skipped" "Skipped scheduled snapshot collection % s

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10303	Information	Protection Set	"Scheduled snapshot skipped" "Skipped scheduled snapshot on volumes associated with ""volume collection %s schedule % s
10304	Information	Protection Set	"Scheduled snapshot succeeded" "Successfully snapshoted volumes associated with ""volume collection %s schedule % s
10305	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the system is unable to log into the vCenter ""server due to an incorrect user name %s or password. Verify the ""user name and password are correct."
10306	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the connection to the vCenter server %s timed ""out. Verify the vCenter server name and IP address."
10307	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because vCenter server %s refused the connection."
10308	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because vCenter server %s is disconnected from the ""Nimble Protection Manager."
10309	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because no route was found to vCenter server %s."
10310	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the socket read from vCenter server %s timed out."
10311	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the network is unreachable."

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10312	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because snapshots are disabled. Some VMWare products
10313	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because %s does not exist."
10314	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because %s does not have permission to perform the ""operation. Configure user with the proper permissions or log in as a ""user with such permissions."
10315	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the system was unable to look up the vCenter ""server %s. Verify the vCenter server name or IP address."
10316	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s due to an encoding error
10317	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the server socket was closed before a reply ""was written."
10318	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because no Datacenter exists."
10319	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because Nimble volume serial number is empty."
10320	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s due to an unexpected error."

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10321	Warning	Protection Set	"Scheduled snapshot succeeded with warning" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the VMFS datastore block size is not large enough ""to accommodate the files in the VM. For VMFS-5 datastore the maximum file ""size is 2TB minus 512 bytes. For the VMFS-3 datastore the maximum file ""size will depend on the block size of the datastore."
10322	Information	Protection Set	"Scheduled snapshot succeeded" "Successfully created snapshot of volumes associated with ""volume collection %s schedule %s"
10323	Information	Protection Set	"Scheduled snapshot succeeded" "Successfully created snapshot of volumes associated with ""volume collection %s schedule % s
10324	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because snapshots are disabled. Some VMware products
10325	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s since the vCenter virtual machine snapshot tasks have not yet completed."
10326	Warning	Protection Set	"Scheduled snapshot failed" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the snapshot name exceeds the maximum allowable ""length on vCenter. Shorten the volume collection name or schedule name."
10327	Information	Protection Set	"Application synchronization failed" "Application synchronization failed for volume collection %s ""schedule % s
10328	Warning	Protection Set	"Scheduled snapshot failed." "Failed to snapshot volumes associated with volume collection % s

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10340	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because %s does not have permission to perform the ""operation. Configure user with the proper permissions or log in as a ""user with such permissions."
10341	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because %s does not exist."
10342	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s due to an encoding error
10343	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because the server socket was closed before a reply ""was written."
10344	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because no Datacenter exists."
10345	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because Nimble volume serial number is empty."
10346	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s due to an unexpected error."
10347	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because the VMFS datastore block size is not large enough ""to accommodate the files in the VM. For VMFS-5 datastore the maximum file ""size is 2TB minus 512 bytes. For the VMFS-3 datastore the maximum file ""size will depend on the block size of the datastore."

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10348	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because snapshots are disabled. Some VMware products
10349	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s since the vCenter virtual machine ""snapshot tasks have not yet completed."
10350	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because the snapshot name exceeds ""the maximum allowable length on vCenter. Shorten the volume collection ""name or schedule name."
10351	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s because virtual machine tools ""in the guest are not running."
10352	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the IP address or hostname %s is not a valid ""vCenter server."
10353	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s. Too many vCenter synchronized snapshots are currently in progress. ""Stagger the schedules for vCenter synchronized snapshots."
10354	Warning	Protection Set	"Scheduled snapshot failed" "Failed to create VSS synchronized snapshot associated with volume collection %s ""schedule %s. Incompatible version of Nimble VSS service on the application server."

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10355	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the system is unable to log into the vCenter ""server due to an incorrect user name %s or password. Verify the ""user name and password are correct."
10356	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the connection to the vCenter server %s timed ""out. Verify the vCenter server name and IP address."
10357	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because vCenter server %s refused the connection."
10358	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because vCenter server %s is disconnected from the ""Nimble Protection Manager."
10359	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because no route was found to vCenter server %s."
10360	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the socket read from vCenter server %s timed out."
10361	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the network is unreachable."
10362	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s because the system was unable to look up the vCenter ""server %s. Verify the vCenter server name or IP address."

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10363	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s due to an un-expected connection error."
10364	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s since the vCenter virtual machine snapshot tasks have not yet completed."
10365	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s due to an un-expected error."
10366	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s due to an unexpected error. ""See vSphere Client to get details on the failure of the snapshot task for this VM."
10367	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s ""schedule %s for virtual machine %s due to failure to quiesce the virtual machine. ""See vSphere Client to get details on the failure of the snapshot task for this VM."
10368	Warning	Protection Set	"Failed to synchronize with VMware vCenter" "Failed to %s vCenter snapshot associated with volume collection %s schedule %s for virtual machine %s due to a retained snapshot. For a given volume collection and schedule
10380	Warning	Protection Set	"Scheduled VSS snapshot failed." "Failed to create snapshot of volumes associated with volume collection %s ""schedule %s because of VSS writer failure. %s"
10381	Warning	Protection Set	"VSS synchronization failed for some of the application objects." "Successfully created snapshot of volumes associated with volume collection %s ""Schedule %s. However

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
10382	Warning	Protection Set	"Scheduled VSS snapshot failed." "Failed to create snapshot of volumes associated with volume collection %s ""schedule %s because of a configuration error. %s"
10383	Warning	Protection Set	"Scheduled VSS snapshot failed." "Failed to create snapshot of volumes associated with volume collection %s ""schedule %s because of an error during snapshot verification process. %s"
10384	Warning	Protection Set	"Scheduled VSS snapshot failed." "Failed to create snapshot of volumes associated with volume collection %s ""schedule %s. %s"
10385	Warning	Protection Set	"Scheduled VSS snapshot failed." "Failed to create snapshot of volumes associated with volume collection %s ""schedule %s. %s"
10386	Warning	Protection Set	"Scheduled VSS snapshot failed." "Failed to create snapshot of volumes associated with volume collection %s ""schedule %s. %s"
12500	Information	iSCSI	"Multi initiator login" "iscsi login to volume %s rejected from %s - existing session from %s"
12503	Information	iSCSI	"Multi initiator login" "iSCSI login to volume %s rejected from %s - existing session from %s"
12057	Error	iSCSI	"iSCSI connections from initiator repeatedly closed by array due to improper target ""IP configuration" "iSCSI connections from initiator %s on volume %s repeatedly closed by array due to ""improper target IP configuration. Review iSCSI connection settings on initiator ""and use the correct target IP. Consider using Nimble Connection Manager to manage ""initiator connections"
12508	Error	iSCSI	"iSCSI connections from initiator repeatedly closed by array due to improper target ""IP configuration" "iSCSI connections with source IP %s and destination IP %s from initiator %s on ""volume %s repeatedly closed by the array. Review iSCSI connection settings on the initiator ""and use the correct target IP."

Volume			
Alert/Event	**Severity**	**Type**	**Description/Possible Cause/Actions**
13101	Warning	Array	"Total number of volumes approaching system limit" "Total number of volumes is %ld which is approaching system limit of %ld"
13102	Information	Array	"Total number of volumes below warning threshold" "Total number of volumes is %ld which is below warning threshold of %ld"
14850	Information	Volume	"Volume fully pinned in cache" "Volume %s fully pinned in cache."
14851	Information	Array	"Pinning of cache pinned volumes may have been affected by a Data Service restart" "Pinning of cache pinned volumes on array %s may have been affected by a Data Service restart. Blocks in any affected cache pinned volumes are being restored to cache."
14852	Information	Array	"Pinning of cache pinned volumes may have been affected by an SSD loss" "Pinning of cache pinned volumes on array %s may have been affected by an SSD loss. Blocks in any affected cache pinned volumes are being restored to cache
14853	Warning	Array	"Pinning of cache pinned volumes may have been affected by an internal error" "Pinning of cache pinned volumes on array %s may have been affected by an internal error. Contact Nimble Storage Support."
14854	Information	Array	"Blocks in any affected cache pinned volumes have been restored to cache" "Blocks in any affected cache pinned volumes on array %s have been restored to cache."
14855	Warning	Array	"Pinning of cache pinned volumes has been affected by an SSD loss" "Pinning of cache pinned volumes on array %s has been affected by an SSD loss. Consider replacing SSD or unpinning some volumes."

Audit Logs

The array OS audit log keeps records of all user-initiated, non-read operations performed on the array. Administrators can view the audit log in a summary table with faceted browsing by time, activity category, and across access type.

Audit Log Messages

Table 23: Audit Active Directory

Active Directory			
Object Type	**Operation Type**	**Event Category**	**Message**
1014	Other	User Access	"Join Domain %s"
1015	Other	User Access	"Leave Domain %s"
1017	Update	User Access	"Edit Domain %s"

Table 24: Audit Alarm

Alarm			
Object Type	**Operation Type**	**Event Category**	**Message**
90017	Update	System Configuration	"Update alarm %s"
90018	Delete	System Configuration	"Delete alarm %s"
90019	Update	System Configuration	"Acknowledge alarm %s"
90020	Update	System Configuration	"Unacknowledge alarm %s"

Table 25: Audit Application Server

Application Server			
Object Type	**Operation Type**	**Event Category**	**Message**
22001	Create	System Configuration	"Create application server record with host %s port %s username %s"
22002	Update	System Configuration	"Update application server record with host %s port %s username %s"
22003	Delete	System Configuration	"Delete application server record with host %s port %s username %s"
22004	Create	System Configuration	"Create application server record %s"
22005	Update	System Configuration	"Update application server record %s"
22006	Delete	System Configuration	"Delete application server record %s"

Table 26: Audit Array

Array			
Object Type	**Operation Type**	**Event Category**	**Message**
14001	Update	System Configuration	"Add array %s to the current group"
14002	Update	System Configuration	"Remove array %s from the current group"
14003	Update	System Configuration	"Update array %s"
14004	Update	System Configuration	"Resetup the current group"
14005	Update	System Configuration	"Complete setup on group %s"
90006	Other	System Configuration	"Halt array %s"
90007	Other	System Configuration	"Halt controller %s on array %s"
90009	Other	System Configuration	"Reboot array %s"
90010	Other	System Configuration	"Reboot controller %s on array %s"
90012	Other	System Configuration	"Failover array %s"

Table 27: Audit Chap User

Chap User			
Object Type	**Operation Type**	**Event Category**	**Message**
4010	Create	Data Access	"Create CHAP user %s"
4011	Update	Data Access	"Update CHAP user %s"
4012	Delete	Data Access	"Delete CHAP user %s"

Table 28: Audit Disk

Disk			
Object Type	**Operation Type**	**Event Category**	**Message**
90002	Update	System Configuration	"Add disk at slot %s of shelf %s on array %s"
90003	Update	System Configuration	"Remove disk from slot %s of shelf %s on array %s"

Table 29: Audit Encrypt Key

Encrypt Key			
Object Type	**Operation Type**	**Event Category**	**Message**
18004	Create	Data Access	"Create master key %s"
18005	Other	Data Access	"Enter passphrase for master key %s"
18006	Update	Data Access	"Clear passphrase for master key %s"
18007	Delete	Data Access	"Delete master key %s"
18008	Update	Data Access	"Change passphrase for master key %s"

Table 30: Audit Audit Fibre Channel Interface

Audit Fibre Channel Interface			
Object Type	**Operation Type**	**Event Category**	**Message**
20001	Update	System Configuration	"Update Fibre Channel interface %s on controller %s of array %s"
20002	Update	System Configuration	"Update Fibre Channel configuration"
20003	Update	System Configuration	"Regenerate Fibre Channel configuration"
20004	Update	System Configuration	"Update Fibre Channel configuration after hardware changes"
20005	Update	System Configuration	"Bulk update Fibre Channel interface(s): %s"

Table 31: Audit Folder

Folder			
Object Type	**Operation Type**	**Event Category**	**Message**
7001	Create	Data Provisioning	"Create folder %s:%s"
7002	Update	Data Provisioning	"Update folder %s:%s"
7003	Delete	Data Provisioning	"Delete folder %s:%s"
7004	Update	Data Provisioning	"Associate volume %s with folder %s:%s"
7005	Update	Data Provisioning	"Dissociate volume %s from folder %s:%s"

Table 32: Audit Initiator

Initiator			
Object Type	**Operation Type**	**Event Category**	**Message**
4015	Create	Data Access	"Add initiator(s) %s to initiator group %s"
4016	Delete	Data Access	"Remove initiator %s from initiator group %s"

Table 33: Audit Initiator Group

Initiator Group			
Object Type	**Operation Type**	**Event Category**	**Message**
4020	Create	Data Access	"Create initiator group %s"
4021	Delete	Data Access	"Delete initiator group %s"
4022	Update	Data Access	"Update initiator group %s"

Table 34: Audit Initiator Group Subnet

Initiator Group Subnet			
Object Type	**Operation Type**	**Event Category**	**Message**
4017	Create	Data Access	"Add subnet(s) %s to initiator group %s"

Initiator Group Subnet			
Object Type	**Operation Type**	**Event Category**	**Message**
4018	Delete	Data Access	"Remove subnet %s from initiator group %s"

Table 35: Audit Internal

Internal			
Object Type	**Operation Type**	**Event Category**	**Message**
1001	Other	User Access	"Login attempt"
1016	Other	User Access	"Elevating user privilege to %s"
12001	Other	Software Update	"Cancel ongoing software download"
12002	Other	Software Update	"Download software update version %s"
12003	Other	Software Update	"Resume software update"
12004	Other	Software Update	"Start software update to version %s"
14006	Other	System Configuration	"Clear group merge state with group %s"
14007	Other	System Configuration	"Validate and merge with group %s"
14008	Update	System Configuration	"Update group configuration"
90001	Update	System Configuration	"Update name for %s %s to %s"
90005	Other	System Configuration	"Initiate AutoSupport"
90008	Other	System Configuration	"Halt group"
90011	Other	System Configuration	"Reboot group"
90014	Other	System Configuration	"Log message"
90015	Update	System Configuration	"Update date to %s%s"
90016	Read	System Configuration	"Validate AutoSupport"
90210	Update	System Configuration	"Set Volume Dedupe for %s to %s"

Table 36: Audit IP

IP			
Object Type	**Operation Type**	**Event Category**	**Message**
16005	Update	System Configuration	"Add IP address %s in network configuration %s"
16006	Update	System Configuration	"Delete IP address %s from network configuration %s"
16007	Update	System Configuration	"Update IP address %s in network configuration %s"

Table 37: Audit Net Config

Net Config			
Object Type	**Operation Type**	**Event Category**	**Message**
16001	Other	System Configuration	"Activate network configuration %s"
16002	Update	System Configuration	"Update network configuration %s"
16003	Delete	System Configuration	"Delete network configuration %s"
16004	Other	System Configuration	"Validate network configuration %s"
16016	Update	System Configuration	"Update iSCSI Connection Method in network configuration %s"

Table 38: Audit NIC

NIC			
Object Type	**Operation Type**	**Event Category**	**Message**
16014	Update	System Configuration	"Assign NIC port %s of array %s to subnet %s in network configuration %s"
16015	Update	System Configuration	"Unassign NIC port %s of array %s from subnet %s in network configuration %s"

Table 39: Audit Partner

Partner			
Object Type	**Operation Type**	**Event Category**	**Message**
10001	Create	Data Protection	"Create partner %s"
10002	Other	Data Protection	"Resume partner %s"
10010	Update	Data Protection	"Update partner %s"
10011	Delete	Data Protection	"Delete partner %s"
10012	Other	Data Protection	"Pause partner %s"

Table 40: Audit Audit Protocol Endpoint

Audit Protocol Endpoint			
Object Type	**Operation Type**	**Event Category**	**Message**
7106	Create	Data Access	"Create protocol endpoint %s"
7107	Update	Data Access	"Update protocol endpoint %s"
7108	Delete	Data Access	"Delete protocol endpoint %s"

Table 41: Audit Audit Performance Policy

Audit Performance Policy			
Object Type	**Operation Type**	**Event Category**	**Message**
2004	Create	Data Provisioning	"Create performance policy %s"
2005	Update	Data Provisioning	"Update performance policy %s"
2006	Delete	Data Provisioning	"Delete performance policy %s"

Table 42: Audit Pool

Pool			
Object Type	**Operation Type**	**Event Category**	**Message**
6001	Create	Data Provisioning	"Create pool %s"
6002	Update	Data Provisioning	"Update pool %s"
6003	Delete	Data Provisioning	"Delete pool %s"
6004	Other	Data Provisioning	"Merge pool %s into pool %s"
6007	Other	Data Provisioning	"Assign array %s to pool %s"
6008	Other	Data Provisioning	"Unassign array %s to pool %s"

Table 43: Audit Audit Protection Policy

Audit Protection Policy			
Object Type	**Operation Type**	**Event Category**	**Message**
8001	Create	Data Protection	"Create volume collection %s"
8002	Update	Data Protection	"Update volume collection %s"
8004	Delete	Data Protection	"Delete volume collection %s"
10006	Other	Data Protection	"Hand over volume collection %s to partner %s"
10007	Other	Data Protection	"Cancel handing over volume collection %s"
10008	Other	Data Protection	"Demote volume collection %s giving ownership to partner %s"
10009	Other	Data Protection	"Promote volume collection %s"
10013	Other	Data Protection	"Validate volume collection %s"

Table 44: Audit Audit Protection Schedule

Audit Protection Schedule			
Object Type	**Operation Type**	**Event Category**	**Message**
8005	Create	Data Protection	"Create protection schedule %s for volume collection %s"
8006	Update	Data Protection	"Update protection schedule %s for volume collection %s"

Audit Protection Schedule			
Object Type	**Operation Type**	**Event Category**	**Message**
8008	Delete	Data Protection	"Delete protection schedule %s for volume collection %s"
8012	Create	Data Protection	"Create protection schedule %s for protection template %s"
8013	Update	Data Protection	"Update protection schedule %s for protection template %s"
8014	Delete	Data Protection	"Delete protection schedule %s for protection template %s"

Table 45: Audit Audit Protection Template

Audit Protection Template			
Object Type	**Operation Type**	**Event Category**	**Message**
8009	Create	Data Protection	"Create protection template %s"
8010	Update	Data Protection	"Update protection template %s"
8011	Delete	Data Protection	"Delete protection template %s"

Table 46: Audit Replication Throttle

Replication Throttle			
Object Type	**Operation Type**	**Event Category**	**Message**
10003	Create	Data Protection	"Create bandwidth throttle for replication partner %s"
10004	Update	Data Protection	"Update bandwidth throttle ID %s for replication partner %s"
10005	Delete	Data Protection	"Delete bandwidth throttle ID %s for replication partner %s"

Table 47: Audit Route

Route			
Object Type	**Operation Type**	**Event Category**	**Message**
16008	Create	System Configuration	"Add route %s in network configuration %s"
16009	Delete	System Configuration	"Delete route %s from network configuration %s"
16010	Update	System Configuration	"Update route %s in network configuration %s"

Table 48: Audit Session

Session			
Object Type	**Operation Type**	**Event Category**	**Message**
1009	Other	User Access	"Kill session"

Table 49: Audit Shelf

Shelf			
Object Type	**Operation Type**	**Event Category**	**Message**
90004	Update	System Configuration	"Add shelf %s on array %s"

Table 50: Audit Snap

Snap			
Object Type	**Operation Type**	**Event Category**	**Message**
4004	Update	Data Protection	"Online snapshot %s of volume %s"
4005	Update	Data Protection	"Offline snapshot %s of volume %s"
8017	Create	Data Protection	"Create snapshot %s of volume %s"
8018	Create	Data Protection	"Create snapshot colletion %s of volumes %s"
8019	Update	Data Protection	"Update snapshot %s of volume %s"
8020	Delete	Data Protection	"Delete snapshot %s of volume %s"
10015	Delete	Data Protection	"Delete replica snapshot %s on downstream partner %s"
10017	Delete	Data Protection	"Delete replica snapshot %s on volume %s for upstream partner %s"

Table 51: Audit Snap Col

Snap Col			
Object Type	**Operation Type**	**Event Category**	**Message**
8021	Create	Data Protection	"Create snapshot collection %s of volume collection %s"
8022	Update	Data Protection	"Update snapshot collection %s of volume collection %s"
8023	Delete	Data Protection	"Delete snapshot collection %s of volume collection %s"

Table 52: Audit SSH Key

SSH Key			
Object Type	**Operation Type**	**Event Category**	**Message**
18001	Create	User Access	"Add SSH key %s for user %s"
18002	Update	User Access	"Update SSH key %s"
18003	Delete	User Access	"Delete SSH key %s for user %s"

Table 53: Audit Subnet

Subnet			
Object Type	**Operation Type**	**Event Category**	**Message**
16011	Create	System Configuration	"Add subnet %s in network configuration %s"
16012	Delete	System Configuration	"Delete subnet %s from network configuration %s"
16013	Update	System Configuration	"Update subnet %s in network configuration %s"

Table 54: Audit User

User			
Object Type	**Operation Type**	**Event Category**	**Message**
1002	Delete	User Access	"Delete user %s"
1003	Update	User Access	"Update user %s"
1004	Other	User Access	"Disable user %s"
1005	Other	User Access	"Enable user %s"
1006	Create	User Access	"Create user %s"
1007	Update	User Access	"Change password for user %s"
1008	Other	User Access	"Logout user %s"
1010	Other	User Access	"User %s session timed out"

Table 55: Audit User Group

User Group			
Object Type	**Operation Type**	**Event Category**	**Message**
1011	Create	User Access	"Create user group %s"
1012	Delete	User Access	"Delete user group %s"
1013	Update	User Access	"Edit user group %s"

Table 56: Audit Volume

Volume			
Object Type	**Operation Type**	**Event Category**	**Message**
2001	Create	Data Provisioning	"Create volume %s"
2002	Update	Data Provisioning	"Update volume %s"
2003	Delete	Data Provisioning	"Delete volume %s"
2007	Create	Data Provisioning	"Clone volume %s from volume %s snapshot %s"
2008	Other	Data Provisioning	"Restore volume %s from snapshot %s"
2009	Create	Data Provisioning	"Create volume %s in folder %s"
4001	Update	Data Provisioning	"Online volume %s"

Volume			
Object Type	**Operation Type**	**Event Category**	**Message**
4002	Update	Data Provisioning	"Offline volume %s"
4003	Update	Data Provisioning	"Online volumes %s"
6005	Other	Data Provisioning	"Move volume %s to pool %s"
6006	Other	Data Provisioning	"Cancel move of volume %s"
6009	Other	Data Provisioning	"Move volume %s to folder %s:%s"
8015	Other	Data Protection	"Associate volume %s with volume collection %s"
8016	Other	Data Protection	"Dissociate volume %s from a previously associated volume collection"
10014	Delete	Data Protection	"Delete replica volume %s on downstream partner %s"
10016	Delete	Data Protection	"Delete replica volume %s for upstream partner %s"

Table 57: Audit Volume ACL

Volume ACL			
Object Type	**Operation Type**	**Event Category**	**Message**
4013	Create	Data Access	"Add ACL for volume %s"
4014	Delete	Data Access	"Remove ACL for volume %s"
4023	Create	Data Access	"Add VMware Virtual Volume ACL for volume %s"
4024	Delete	Data Access	"Remove VMware Virtual Volume ACL for volume %s"
7109	Create	Data Access	"Add ACL record for protocol endpoint %s"
7110	Delete	Data Access	"Remove ACL record for protocol endpoint %s"

System and Timeout Limits

This section summarizes limits for systems (both iSCSI and Fibre Channel arrays) and timeout values.

System Limits

This section summarizes the system limits for HPE Nimble Storage arrays. Alerts or log messages are generated when the system reaches the warning count, so this is a *threshold*. No more objects of the specified type can be created when the system reaches the maximum count, so this is a *limit*.

Important: Volume scalability to 10,000 volumes per group is supported when the group leader is an array model that supports 10,000 volumes, and it belongs to a pool that can also support this limit. Pool limits are the same as the limits of the smallest capacity array in the pool.

See the notes at the end of the table for information about array models that support specific limits.

See Relationship of Groups, Pools, Arrays, Folders, and Volumes on page 16 for information about the scope of objects in single-array and multi-array groups.

Table 58: System Limits

Object Type	Scope	iSCSI Maximum	FC Maximum	Warning	Notes
Array	Group	4	4	4	
	Pool	4	4	4	
Branch	Group	10,000	10,000	9,500	
CHAP user	Group	1024	N/A	960	
Data connection	Controller	80000 15000	80000 15000	72000 13500	An iSCSI or a Fibre Channel path to a volume 1 (80000) 2 (15000)
Session	Array	12000	12000	None	An iSCSI session or FC login
FC initiator	FC port	N/A	256	None	Prior to NimbleOS 4.x, this value was 128.
Folder	Pool	128	128	None	
	Group	128	128	None	
Initiator	Initiator group	256	256	240	
	Group	10,000	10,000	9,500	
Initiator group	Subnet	60	N/A	54	
	Group	1,024	1,024	960	

Object Type	Scope	iSCSI Maximum	FC Maximum	Warning	Notes
Network configuration	Array	4	4	4	
Performance policy	Group	100	100	90	
Pool	Group	4	4	4	
Protection schedule	Group	5,000	5,000	4,500	
	Protection template	10	10	8	
	Volume collection	10	10	8	
Protection template	Group	50	50	45	
Replication bandwidth policy (throttle)	Group	50	50	45	
	Replication partner	25	25	20	
Replication partner	Group	50	50	45	
	Pool	50	50	45	
Route	Network configuration	10	10	8	
SSH key	User	10	10	8	
Snapshot	Group	300,000 190,000	300,000 190,000	275,000 175,000	1 2 3
	Pool	300,000 190,000	300,000 190,000	275,000 175,000	1 2 3
	Volume	1,000	1,000	900	
Writable Snapshot	Array	40000 10000	40000 10000	None	1 2 4
Snapshot Collection	Volume Collection	1,000	1,000	900	
Subnet	Network configuration	60	60	54	
User account	Group	100	100	90	

Object Type	Scope	iSCSI Maximum	FC Maximum	Warning	Notes
Volume	Group	10,000	10,000	9,200	1
		1,024	1,024	960	2
	Pool	10,000	10,000	9,200	1
		1,024	1,024	960	2
	Performance policy	10,000	10,000	9,200	1
		1,024	1,024	960	2
	Volume collection	50	50	45	Volumes in a collection can exist in different pools.
	Volume limit	127 TiB	127 TiB	N/A	
Volume access control list (ACL)	Group	65,536	65,536	60,000	
	Pool	65,536	65,536	60,000	
	Volume	64	64	60	
Volume collection	Group	2,000	2,000	1,800	1
		512	512	480	2

Note:

1 Limits apply to AF5000, AF7000, AF9000, AF60, and AF80 model arrays

2 Limits apply to all model arrays not listed in Note 1.

3 A frequent schedule is a schedule in a volume collection that triggers snapshots more often than every five minutes. There is a limitation of five frequent schedules per array group. There is no restriction on schedules with a period of five minutes or more.

4 Writable snapshots also count against the total Snapshot limit.

Timeout Values

Timeout Values for iSCSI and FC Arrays

Various timeout values affect storage array targets from Windows hosts with and without Multipath-IO (MPIO). Refer to these recommended values for various configurations. Note that the *LinkDownTime* and *MaxRequestHoldTime* values apply to iSCSI arrays only. The HPE Storage Toolkit for Windows sets all these values.

Table 59: Timeout Values

Registry Key	Meaning of Value	Default Values (in seconds)	
		Windows	Storage Array
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\service\Disk\TimeOutValue	Global disk timeout value for disk.sys. If timeout value is met, OS sends bus reset to all LUNs.	60 or 120 (OS dependent)	60
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\Control\Class\ {4D36E97B-E325-11CE-BFC1-08002BE10318} \0001 \Parameters\MaxRequestHoldTime (See Note)	(iSCSI only) Maximum time (in seconds) for which requests will be queued if connection to the target is lost and the connection is being retried. (Applies only to non-MPIO disks.)	60	120
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\Control\Class\ {4D36E97B-E325-11CE-BFC1-08002BE10318} \0001\Parameters\LinkDownTime (See Note)	(iSCSI only) Determines how long (in seconds) I/O requests will be held in the device queue for a MPIO path and retried if the connection to the target is lost.	15	30
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\services\mpio\Parameters \PDORemovePeriod	Time (in seconds) that the multipath pseudo-LUN (Physical Device Object in the kernel) will continue to remain in system memory, even after losing all paths to the volume.	20	90
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\services\mpio\Parameters \UseCustomPathRecoveryInterval	*UseCustomPathRecoveryInterval* (Boolean 1 or 0) and *PathRecoveryInterval* times (in seconds) can be used to perform path recovery while the *PDORemovalPeriod* timer counts down.	0	1
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\services\mpio\Parameters \PathRecoveryInterval		40	40
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\services\mpio\Parameters \RetryCount	Number of times a failed I/O if DSM determines that a failing request must be retried on current path to a LUN on certain transient errors.	3	30
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\services\mpio\Parameters \RetryInterval	Interval of time (in seconds) after which a failed request is retried (determined by DSM, and assumes that the I/O has been retried a fewer number of times than *RetryCount*).	1	1
HKEY_LOCAL_MACHINE\SYSTEM \CurrentControlSet\services\msdsm\Parameters \DsmMaximumStateTransitionTime	Time SCSI Reservation is retried on a failed path for specific CHECK_CONDITIONS raised by the array.	3	120

Note:

- All values are in seconds.
- A reboot is required after value changes.
- The iSCSI initiator driver instance ID (for iSCSI only values) may be different than 0001.

FC HBA Timeout Values

The following values apply when configuring timeouts on FC HBA ports, depending on the type. The Emulex management tool *OneCommand Manager* allows you to configure timeouts on a per-port basis. Similarly, Qlogic's *QconvergeConsole* allows per-port configuration under **Port** > **Parameters** > **HBA Parameters/Advanced HBA Parameters**. Note that reboots are required for Qlogic FC HBAs.

Table 60: FC HBA Timeout Values

Timeout Type	Value	Meaning of Value	Default Value (in seconds)	Range (in seconds)
Emulex FC HBA	LinkTimeout	Time interval after which a link that is down stops issuing a BUSY status for requests and starts issuing SELECTION_TIMEOUT error status. This includes port login and discovery time. (No reboot required.)	30	0-255
	NodeTimeout	Time interval after which a prior logged-in node issues SELECTION_TIMEOUT error status to an I/O request. This causes the system to wait for a node that might re-enter the configuration soon before reporting a failure. The timer starts after port discovery is completed and the node is no longer present. (No reboot required.)	30*	0-255
Qlogic FC HBA	LinkDownTimeout	Timeout value when a link that is down stops issuing a BUSY status for requests and starts issuing SELECTION_TIMEOUT error status. This LinkDownTimeOut includes port login and discovery time (Reboot required.)	30	0-255
	PortDownRetryCount	Number of times the I/O request is resent to a port that is not responding in one second intervals. Alternately, this setting specifies the number of seconds the software waits to retry a command to a port returning port down status. (Reboot required.)	30	0-255

Note: If you need to manage unused devices, you can perform a cleanup to remove them from the registry. You can also use the DevNodeClean utility from the Microsoft Download Center.

Important: *By default, an Emulex HBA sets the NodeTimeout value to 30 seconds. In the event of a path failure, this can cause I/O to pause on the host for the assigned timeout. We recommend that you set this value to 1, which allows the host to immediately resume I/O.

Firewall Ports

Configure local ports for HTTP, HTTPS, iSCSI, SNMP, SSH (incoming and outgoing), and other data and management protocols.

Configure Firewall Ports

Use this information to configure local ports for incoming and outgoing HTTP, HTTPS, iSCSI, SCP, SNMP, SSH, TCP, and other data and management protocols.

Table 61: Group Egress Ports – Nimble External Port

Port Number	Service	Protocol	Destination DNS/IP
443 TCP	DNA, heartbeat	HTTPS	nsdiag.nimblestorage.com
443 TCP	Storage array alerts *	HTTPS	nsalerts.nimblestorage.com
443 TCP	Storage array statistics	HTTPS	nsstats.nimblestorage.com
443 TCP	Software downloads	HTTPS	update.nimblestorage.com
2222 TCP	Secure tunnel	SSH	hogan.nimblestorage.com
4311 TCP	Nimble Protection Manager	SOAP/HTTP	application server IP **
8443 TCP	vCenter VASA/VVol integration	HTTPS	Management IP address and both diagnostic IP addresses

* An array sends DNA messages using HTTPS POST back to Nimble Storage Support, if it is enabled. If three HTTPS POST attempts are made and they all fail, these notifications will revert to email relay.

** If the application server connecting with these ports on a Nimble array is on the same side of the firewall as the array, you do not need to open these ports in the firewall.

Note: The array may initiate connections to these external addresses from the Management IP or any controller support IP address.

Note: When configuring firewall rules for the destinations listed above, it is recommended that you specify the destination by host name rather than by IP address, and allow DNS to resolve the IP address. In the event that there is a change in the publicly available IP address for one of these destinations, the change will be communicated by a notification on the InfoSight portal. Other methods of sending notifications of such changes may be chosen as needed.

Table 62: Intra-group Ports – TCP Ports Needed Between Arrays in a Group

Port Number	Service	Protocol	IP Address
4211 TCP	Array setup (incoming) and management (intra-group)	SOAP/HTTP	Data IP(s)
4212 TCP	Group controller management	HTTP	Data IP(s)
4241 TCP	Group controller management	DTS	Data IP(s)
5432 TCP	Group configuration synchronization	DTS	Data IP(s)

Port Number	Service	Protocol	IP Address
5521 TCP	Group data services	DTS	Data IP(s)
5706 TCP	Group event reporting	SOAP/HTTP	Data IP(s)
6716 TCP	DSD miscellaneous management	SOAP/HTTP	Data IP(s)
6717 TCP	GMD array management (GAI)	SOAP/HTTP	Data IP(s)
6718 TCP	Group controller management	DTS	Data IP(s)
6719 TCP	Data forwarding	DTS	Data IP(s)
6720 TCP	Bin migration	DTS	Data IP(s)
6721 TCP	Bin map management – DSD	DTS	Data IP(s)
6722 TCP	iSCSI	DTS	Data IP(s)
6723 TCP	Bin map management - GDD	DTS	Data IP(s)
6724 TCP	iSCSI	DTS	Data IP(s)
6725 TCP	DSD volume management	SOAP/HTTP	Data IP(s)
6726 TCP	SCSI	DTS	Data IP(s)
6727 TCP	SCSI	DTS	Data IP(s)
6728 TCP	Key Protocol	DTS	Data IP(s)
6729 TCP	LU cache (DSD-GDD)	DTS	Data IP(s)
6730 TCP	Key Protocol	DTS	Data IP(s)
6731 TCP	LU cache (DSD-DSD)	DTS	Data IP(s)

Note: If the arrays within the group are on the same side of the firewall, you do not need to open these ports in the firewall.

Table 63: Inter-group Ports – TCP Ports Needed Between Replication Partners

Port Number	Service	Protocol	IP Address
4213 TCP	Replication control (exchange of replication configuration information between groups)	SOAP/HTTP	Management IP address and both diagnostic IP addresses of all replication partners and group members
4214 TCP	Replication data (transfer of replicated data)	NS-REPL	**Note:** Use either: 1 — All IP addresses in the management subnet of all replication partners and group members or 2 — All data IP addresses in the chosen data subnet of all replication partners and group members *

Port Number	Service	Protocol	IP Address
5391 TCP	Secure web-service communications Exchange of SSL keys for encrypted volumes	SOAP/HTTPS	Management IP address and both diagnostic IP addresses

* Assumes that all replication partners were chosen to perform replication transfer over the data subnet.

Important:

There are two options for replication:

1. Replication transfer and replication over the Management subnet.
2. Replication transfer over data subnet specified during replication partner configuration. Replication control is still transferred over the management subnet per table.

Note: If the arrays in the two groups are on the same side of the firewall, you do not need to open these ports in the firewall.

Table 64: Group Ingress Ports – Nimble External Ports

Port Number	Service	Protocol	IP Address
22 TCP	Group management (CLI)	SSH	Management IP address and both diagnostic IP addresses
161 UDP	SNMP get	SNMP	Management IP address and both diagnostic IP addresses
redirect 80 TCP to 443 TCP	Group management (GUI), redirects to 443 TCP	HTTP	Management IP address and both diagnostic IP addresses
443 TCP / 5392 TCP	Group management (GUI)	HTTPS	Management IP address and both diagnostic IP addresses
3260 TCP	SNMP statistics	iSCSI	Data IP(s) and discovery IP(s)
4210 TCP	Group management (GUI charts and NPM)	SOAP/HTTP	Management IP address and both diagnostic IP addresses
4211 TCP	Array setup (incoming) and management (intra-group)	SOAP/HTTP	Data IP(s)
5988 TCP	CIM server **	HTTP	Management IP address and both diagnostic IP addresses
5989 TCP	CIM server	HTTPS/CIM-XML	Management IP address and both diagnostic IP addresses
5390 TCP	Secure web-service communications	SOAP/HTTPS	Data IP(s)
5391 TCP *	Third-party agents and utilities	SOAP/HTTPS	Management IP address and both diagnostic IP addresses *
5392 TCP *	Group management, third-party agents and utilities	Nimble REST API	Management IP address and both diagnostic IP addresses *

Port Number	Service	Protocol	IP Address
5393 TCP	Array Management, third party utilities and agents	HTTPS	Management IP address and both diagnostic IP addresses *
8443 TCP	vCenter VASA/VVol integration	HTTPS	Management IP address and both diagnostic IP addresses

* Some third-party utilities may use both TCP port 5391 and TCP port 5392. Refer to the relevant integration guides available on InfoSight, or from the third-party software vendor for more information.

** Fibre Channel arrays do not use the CIM server (cimserver) service, so port 5989 does not need to be open on them.

Note: If the client and the Nimble arrays within the group are on the same side of the firewall, you do not need to open these ports in the firewall.

Table 65: Group Egress Ports – Other External Ports

Port Number	Service	Protocol	Destination DNS/IP
25 * UDP & 25 * TCP	SMTP	SMTP	SMTP server IP
53 / UDP & 53 TCP	DNS	DNS	DNS server IP
123 / UDP	NTP	NTP	NTP server IP
162 * / UDP	SNMP trap	SNMP	SNMP trap listener
443 TCP	HTTPS	HTTPS	vCenter IP
514 UDP	Syslogd	UDP	Syslog server IP
4311 TCP	Microsoft VSS	VSS	Application server IP
Configurable TCP	HTTP	HTTP	HTTP proxy server IP
53 TCP/UDP ** 88 TCP/UDP ** 123 UDP *** 137 TCP/UDP 139 TCP/UDP 389 TCP/UDP	Active Directory Authentication	DNS, Kerberos, SMB	All Active Directory domain controllers

* Default, but can be changed.

** DNS services should be provided by the domain controller, or by an alternative with the appropriate zones and AD records.

*** Array should be configured to use the Active Directory server as the NTP server, or the array and domain controllers should be configured to use the same NTP server. Array clock must remain within 5 minutes of the domain controller clock, or domain authentication will fail.

Note: If the service is on the same side of the firewall as the array, you do not need to open these ports in the firewall.

Regulatory and Safety Information

For important safety, environmental, and regulatory information, refer to this section and to the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products* guide available at http://www.hpe.com/support/Safety-Compliance-EnterpriseProducts.

Regulatory Warnings

European Union (EU), Australia, New Zealand

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

Israel

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

אזהרה

מוצר זה הוא מוצר Class A.
בסביבה ביתית,מוצר זה עלול לגרום הפרעות בתדר רדיו,ובמקרה זה ,המשתמש עשוי להידרש לנקוט אמצעים מתאימים.

Korea

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

경고:
이 기기는 업무용(A급) 전자파적합기기로서 판매자 또는 사용자는 이 점을 주의하시기 바라며, 가정외의 지역에서 사용하는 것을 목적으로 합니다.

Taiwan

Warning

This is a Class A product. In a domestic environment this product may cause radio interference in which case the user may be required to take adequate measures.

警告:
警告使用者：這是甲類的資訊產品，在 居住的環境中使用時，可能會造成射頻干擾，在這種情況下，使用者會被要求 採取某 些適當的對策

Battery Safety

For those Nimble arrays equipped with a battery to power the NVRAM, those batteries are not user-replaceable. However, the following safety precautions are mandated.

CAUTION:

- There is a risk of explosion if battery is replaced by an incorrect type. Dispose of used batteries according to the instructions.
- Battery may explode if mistreated. Do not disassemble or dispose of in fire. Dispose of used batteries promptly. Keep away from children.