



NetApp Element 12.0

User Guide

June 2021 | 215-14901_2021-06_en-us
doccomments@netapp.com

Contents

About this guide.....	6
SolidFire storage system	7
Clusters.....	8
Nodes.....	8
Storage nodes.....	9
Fibre Channel nodes.....	9
Drives.....	9
Custom protection domains.....	10
Management node for Element software.....	10
Management services for SolidFire all-flash storage.....	10
Persistent volumes.....	10
SolidFire Active IQ.....	11
SolidFire software interfaces.....	11
Networking.....	12
Switch configuration for clusters running Element software.....	13
Network port requirements.....	13
System setup.....	14
Setup overview.....	14
Determining which SolidFire components to install.....	15
Setting up an Element storage system.....	15
Configuring a storage node.....	16
Creating a storage cluster.....	19
Accessing the Element software user interface.....	20
Adding drives to a cluster.....	20
Configuring a Fibre Channel node.....	21
Setting up a management node.....	23
Configuring SolidFire system options after deployment.....	24
Changing the Element software default SSL certificate.....	24
Upgrading storage nodes	25
Updating management services for Element storage-based installations.....	26
Using basic options in the Element software UI.....	27
Accessing the Element software user interface.....	27
Limiting results by using filters.....	28
Sorting lists.....	28
Viewing API activity.....	28

Interface refresh rate impacted by cluster load	29
Icons in the Element interface.....	29
Providing feedback.....	30
System management.....	31
Managing cluster administrator user accounts.....	31
Storage cluster administrator account types.....	32
Cluster Admins details.....	32
Creating a cluster administrator account.....	33
Editing cluster administrator permissions.....	34
Changing passwords for cluster administrator accounts.....	34
Managing LDAP.....	34
Enabling Multi-factor authentication.....	41
Setting up Multi-factor authentication.....	42
Additional information for Multi-factor authentication.....	42
Configuring cluster settings.....	43
Setting cluster full threshold.....	44
Enabling and disabling support access.....	44
Enabling and disabling encryption for a cluster.....	44
Encryption at rest.....	45
Managing the Terms of Use banner.....	45
Enabling a broadcast client.....	46
Managing SNMP.....	46
Managing drives.....	48
Managing nodes.....	49
Viewing Fibre Channel ports details.....	53
Fibre Channel ports details.....	53
Managing virtual networks.....	54
Creating a cluster supporting FIPS drives.....	57
Avoiding mixing nodes for FIPS drives.....	58
Enabling encryption at rest.....	58
Identifying whether nodes are ready for the FIPS drives feature.....	58
Enabling the FIPS drives feature.....	59
Checking the FIPS drive status.....	59
Troubleshooting the FIPS drive feature.....	59
Enabling FIPS 140-2 for HTTPS on your cluster.....	60
SSL ciphers.....	61
Getting started with External key management.....	62
Setting up External key management.....	62
Recovering inaccessible or invalid Authentication Keys.....	63
External Key Management API Commands.....	63
Data management.....	65
Working with user accounts.....	65
Creating an account.....	65
Account details.....	66
Viewing individual account performance details	66
Editing an account.....	67
Deleting an account.....	67
Working with volumes.....	68
Quality of Service.....	69
QoS policies.....	71

Creating a volume.....	73
Volume details.....	74
Viewing individual volume performance details.....	75
Editing active volumes.....	75
Deleting a volume.....	76
Restoring a deleted volume.....	77
Purging a volume.....	77
Cloning a volume.....	77
Assigning LUNs to Fibre Channel volumes.....	79
Applying a QoS policy to volumes.....	79
Removing the QoS policy association of a volume.....	79
Working with virtual volumes.....	80
Enabling virtual volumes.....	81
Viewing virtual volume details.....	82
Virtual volume details.....	82
Individual virtual volume details.....	83
Deleting a virtual volume.....	84
Storage containers.....	85
Protocol endpoints.....	87
Bindings.....	88
Host details.....	88
Working with volume access groups and initiators.....	89
Creating a volume access group.....	90
Volume access group details.....	91
Viewing individual access group details.....	92
Adding volumes to an access group.....	92
Removing volumes from an access group.....	92
Creating an initiator.....	93
Editing an initiator.....	93
Adding a single initiator to a volume access group.....	94
Adding multiple initiators to a volume access group.....	95
Removing initiators from an access group.....	95
Deleting an access group.....	96
Deleting an initiator.....	96
Data protection.....	97
Using volume snapshots for data protection.....	97
Using individual volume snapshots for data protection task.....	98
Using group snapshots for data protection task.....	102
Scheduling a snapshot.....	107
Performing remote replication between clusters running NetApp Element software.....	110
Planning cluster and volume pairing for real-time replication.....	111
Pairing clusters.....	111
Pairing volumes.....	114
Validating volume replication.....	119
Deleting a volume relationship after replication.....	119
Managing volume relationships.....	120
Using SnapMirror replication between Element and ONTAP clusters.....	124
SnapMirror overview.....	124
Enabling SnapMirror on the cluster.....	124
Enabling SnapMirror on the volume.....	125
Creating an endpoint.....	125
Creating a SnapMirror relationship.....	126

SnapMirror relationship actions.....	127
SnapMirror labels.....	128
Disaster recovery using SnapMirror.....	129
Backing up and restoring volumes.....	135
Backing up a volume to an Amazon S3 object store.....	135
Backing up a volume to an OpenStack Swift object store.....	136
Backing up a volume to a SolidFire storage cluster.....	136
Restoring a volume from backup on an Amazon S3 object store.....	137
Restoring a volume from backup on an OpenStack Swift object store.....	137
Restoring a volume from backup on a SolidFire storage cluster.....	138
System monitoring and troubleshooting	140
Viewing information about system events.....	141
Event types.....	141
Viewing status of running tasks.....	143
Viewing system alerts.....	144
Cluster fault codes.....	145
Viewing node performance activity.....	155
Viewing volume performance.....	156
Volume performance details.....	156
Viewing iSCSI sessions.....	157
iSCSI session details.....	157
Viewing Fibre Channel sessions.....	158
Fibre Channel session details.....	158
Troubleshooting drives.....	159
Removing failed drives from the cluster.....	160
Basic MDSS drive troubleshooting.....	160
Adding MDSS drives.....	161
Removing MDSS drives.....	162
Troubleshooting nodes.....	162
Powering down a cluster.....	162
Working with per-node utilities for storage nodes.....	163
Accessing per-node settings using the per-node UI.....	163
Network settings details from the per-node UI.....	164
Cluster settings details from the per-node UI.....	165
Running system tests using the per-node UI.....	166
Running system utilities using the per-node UI.....	167
Working with the management node.....	169
Understanding cluster fullness levels.....	169
Contacting NetApp Support.....	171
Where to find product documentation and other information	172
Copyright and trademark.....	173
Copyright.....	173
Trademark.....	173

About this guide

The *NetApp Element Software User Guide* provides information about how to configure, manage, and use storage systems running Element data management software. This guide is intended for IT professionals, software developers, and others who install, administer, or troubleshoot storage systems running NetApp Element software.

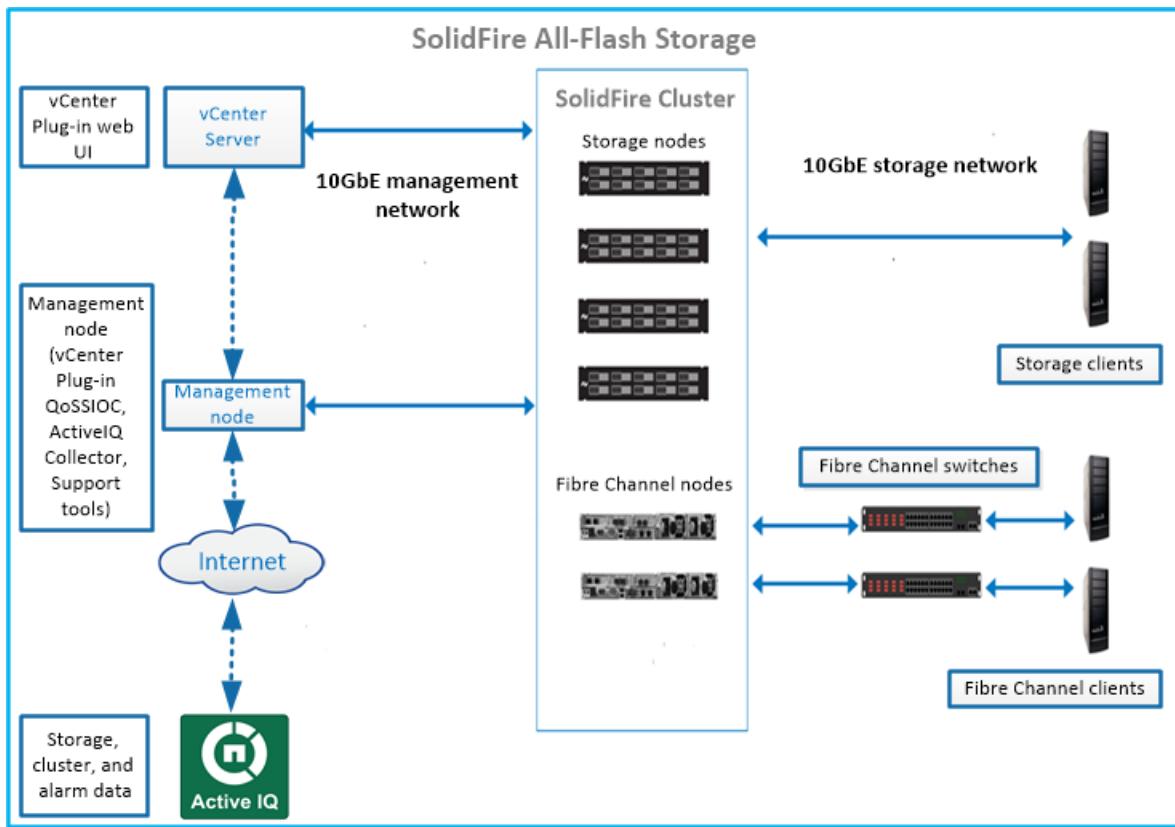
This guide makes the following assumptions:

- You have a background as a Linux system administrator.
- You are familiar with server networking and networked storage, including IP addresses, netmasks, and gateways.

SolidFire storage system

A SolidFire all-flash storage system is comprised of discrete hardware components (drive and nodes) that are combined into a single pool of storage resources through NetApp Element software running independently on each node. This unified cluster presents as a single storage system for use by external clients and is managed as a single entity through the Element software UI, API and other management tools.

Using the NetApp Element software user interface, you can set up and monitor SolidFire cluster storage capacity and performance, and manage storage activity across a multi-tenant infrastructure.



A SolidFire all-flash storage system includes the following components:

- Nodes: Physical hardware that provides the storage resources for a cluster. There are two types of nodes:
 - Storage nodes: Servers containing a collection of drives.
 - Fibre Channel (FC) nodes: Used to connect FC clients via a Fibre Channel switch.
- Cluster: The hub of the SolidFire storage system comprised of a least four nodes.
- Management node: Enables you to upgrade and provide system services including monitoring and telemetry, manage cluster assets and settings, run system tests and utilities, and provide NetApp Support access for troubleshooting. The management node (mNode) is a virtual machine that runs in tandem with an Element software-based storage cluster.
- Active IQ: A web-based tool that provides continually updated historical views of cluster-wide data. You can set up alerts for specific events, thresholds, or metrics. Active IQ enables you to monitor system performance and capacity, as well as stay informed about cluster health.

- Drives are used in storage nodes and store data for the cluster. A storage node contains two types of drives:
 - Volume metadata drives store information that defines the volumes and other objects within a cluster.
 - Block drives store data blocks for application volumes.

Clusters

A cluster is the hub of a SolidFire storage system and is made up of a collection of nodes. You must have at least four nodes in a cluster for SolidFire storage efficiencies to be realized. A cluster appears on the network as a single logical group and can then be accessed as block storage.

Creating a new cluster initializes a node as communications owner for a cluster and establishes network communications for each node in the cluster. This process is performed only once for each new cluster. You can create a cluster by using the Element UI or the API.

You can scale out a cluster by adding additional nodes. When you add a new node, there is no interruption of service and the cluster automatically uses the performance and capacity of the new node.

Administrators and hosts can access the cluster using virtual IP addresses. Any node in the cluster can host the virtual IP addresses. The management virtual IP (MVIP) enables cluster management through a 1GbE connection, while the storage virtual IP (SVIP) enables host access to storage through a 10GbE connection. These virtual IP addresses enable consistent connections regardless of the size or makeup of a SolidFire cluster. If a node hosting a virtual IP address fails, another node in the cluster begins hosting the virtual IP address.

Note: Beginning in Element version 11.0, nodes can be configured with IPv4, IPv6, or both addresses for their management network. This applies to both storage nodes and management nodes, except for management node 11.3 and later which does not support IPv6. When creating a cluster, only a single IPv4 or IPv6 address can be used for the MVIP and the corresponding address type must be configured on all nodes.

Nodes

Nodes are the individual hardware components that are grouped into a cluster to be accessed as block storage. There are two types of nodes in a SolidFire storage system: storage nodes and Fibre Channel nodes.

Related concepts

[Storage nodes](#) on page 9

A SolidFire storage node is a server containing a collection of drives that communicate with each other through the Bond10G network interface. Drives in the node contain block and metadata space for data storage and data management.

[Fibre Channel nodes](#) on page 9

SolidFire Fibre Channel nodes provide connectivity to a Fibre Channel switch, which you can connect to Fibre Channel clients. Fibre Channel nodes act as a protocol converter between the Fibre Channel and iSCSI protocols; this enables you to add Fibre Channel connectivity to any new or existing SolidFire cluster.

[Drives](#) on page 9

A storage node contains one or more physical drives that are used to store a portion of the data for the cluster. The cluster utilizes the capacity and performance of the drive after the drive has been successfully added to a cluster.

Storage nodes

A SolidFire storage node is a server containing a collection of drives that communicate with each other through the Bond10G network interface. Drives in the node contain block and metadata space for data storage and data management.

Storage nodes have the following characteristics:

- Each node has a unique name. If a node name is not specified by an administrator, it defaults to SF-XXXX, where XXXX is four random characters generated by the system.
- Each node has its own high-performance non-volatile random access memory (NVRAM) write cache to improve overall system performance and reduce write latency.
- Each node is connected to two networks, storage and management, each with two independent links for redundancy and performance. Each node requires an IP address on each network.
- You can create a cluster with new storage nodes, or add storage nodes to an existing cluster to increase storage capacity and performance.
- You can add or remove nodes from the cluster at any time without interrupting service.

Fibre Channel nodes

SolidFire Fibre Channel nodes provide connectivity to a Fibre Channel switch, which you can connect to Fibre Channel clients. Fibre Channel nodes act as a protocol converter between the Fibre Channel and iSCSI protocols; this enables you to add Fibre Channel connectivity to any new or existing SolidFire cluster.

Fibre Channel nodes have the following characteristics:

- Fibre Channel switches manage the state of the fabric, providing optimized interconnections.
- The traffic between two ports flows through the switches only; it is not transmitted to any other port.
- Failure of a port is isolated and does not affect operation of other ports.
- Multiple pairs of ports can communicate simultaneously in a fabric.

Drives

A storage node contains one or more physical drives that are used to store a portion of the data for the cluster. The cluster utilizes the capacity and performance of the drive after the drive has been successfully added to a cluster.

A storage node contains two types of drives:

Volume metadata drives

These drives store compressed information that defines each volume, clone, or snapshot within a cluster. The total metadata drive capacity in the system determines the maximum amount of storage that can be provisioned as volumes. The maximum amount of storage that can be provisioned is independent from how much data is actually stored on the block drives of the cluster. Volume metadata drives store data redundantly across a cluster using Double Helix data protection.

Note: Some system event log and error messages refer to volume metadata drives as slice drives.

Block drives

These drives store the compressed, de-duplicated data blocks for server application volumes. Block drives make up a majority of the storage capacity of the system. The

majority of read requests for data already stored on the SolidFire cluster, as well as requests to write data, occur on the block drives. The total block drive capacity in the system determines the maximum amount of data that can be stored, taking into account the effects of compression, thin provisioning, and de-duplication.

Custom protection domains

You can define a custom protection domain layout, where each node is associated with one and only one custom protection domain. By default, each node is assigned to the same default custom protection domain.

If no custom protection domains are assigned:

- Cluster operation is unaffected.
- Custom level is neither tolerant nor resilient.

If more than one custom protection domain is assigned, each subsystem will assign duplicates to separate custom protection domains. If this is not possible, it reverts to assigning duplicates to separate nodes. Each subsystem (for example, bins, slices, protocol endpoint providers, and ensemble) does this independently.

Note: Using custom protection domains assumes that no nodes share a chassis.

The following API methods expose these new protection domains:

- `GetProtectionDomainLayout` - shows which chassis and which custom protection domain each node is in.
- `SetProtectionDomainLayout` - allows a custom protection domain to be assigned to each node.

Contact NetApp support for further details on using custom protection domains.

Related information

[Managing storage with the Element API](#)

Management node for Element software

The management node (mNode) is a virtual machine that runs in parallel with one or more Element software-based storage clusters. It is used to upgrade and provide system services including monitoring and telemetry, manage cluster assets and settings, run system tests and utilities, and enable NetApp Support access for troubleshooting.

As of the Element 11.3 release, the management node functions as a microservice host, allowing for quicker updates of select software services outside of major releases. These microservices or management services, such as the Active IQ collector, QoSIOC for the vCenter Plug-in, and mNode service, are updated frequently as service bundles. Additional services including HealthTools for storage node software upgrades and support tools (remote support tunneling) are also available from the management node.

Management services for SolidFire all-flash storage

Management services provide central and extended management functionality for SolidFire all-flash storage. These services include Active IQ system telemetry, logging, and service updates, as well as the QoSIOC service for the Element Plug-in for vCenter.

Persistent volumes

Persistent volumes allow management node configuration data to be stored on a specified storage cluster, rather than locally with a VM, so that data can be preserved in the event of management

node loss or removal. Persistent volumes are an optional but recommended management node configuration.

An option to enable persistent volumes is included in the installation and upgrade scripts when deploying a new management node. Persistent volumes are volumes on an Element software-based storage cluster that contain management node configuration information for the host management node VM that persists beyond the life of the VM. If the management node is lost, a replacement management node VM can reconnect to and recover configuration data for the lost VM.

Persistent volumes functionality, if enabled during installation or upgrade, automatically creates multiple volumes with NetApp-HCI- pre-pended to the name on the assigned cluster. These volumes, like any Element software-based volume, can be viewed using the Element software web UI, NetApp Element Plug-in for vCenter Server, or API, depending on your preference and installation. Persistent volumes must be up and running with an iSCSI connection to the management node to maintain current configuration data that can be used for recovery.



Attention: Persistent volumes that are associated with management services are created and assigned to a new account during installation or upgrade. If you are using persistent volumes, do not modify or delete the volumes or their associated account.

SolidFire Active IQ

Active IQ is a web-based tool that provides continually updated historical views of cluster-wide data. You can set up alerts for specific events, thresholds, or metrics. Active IQ enables you to monitor system performance and capacity, as well as stay informed about cluster health.

You can find the following information about your system in Active IQ:

- Number of nodes and status of the nodes: healthy, offline, or fault
- Graphical representation of CPU and memory usage
- Details about the node, such as serial number, slot location in the chassis, model, and version of NetApp Element software running on the storage node
- CPU and storage-related information about the virtual machines

SolidFire software interfaces

You can manage a SolidFire storage system by using NetApp Element software interfaces and integration utilities.

NetApp Element software user interface

Enables you to set up SolidFire storage, monitor cluster capacity and performance, and manage storage activity across a multi-tenant infrastructure. Element is the storage operating system at the heart of a SolidFire cluster. Element software runs independently on all nodes in the cluster and enables the nodes of the cluster to combine resources and present as a single storage system to external clients. Element software is responsible for all cluster coordination, scale and management of the system as a whole. The software interface is built upon the Element API.

NetApp Element Plug-in for vCenter Server

Enables you to configure and manage storage clusters running Element software. The plug-in provides an alternative interface for the Element UI within VMware vSphere.

NetApp Element software API

Enables you to use a set of objects, methods, and routines to manage SolidFire storage. The Element API is based on the JSON-RPC protocol over HTTPS. You can monitor API

operations in the Element UI by enabling the API Log; this enables you to see the methods that are being issued to the system. You can enable both requests and responses to see how the system replies to the methods that are issued.

Management node UIs

The management node contains two UIs: a UI for managing REST-based services and a per-node UI for managing network and cluster settings and operating system tests and utilities. From the REST API UI, you can access a menu of service-related APIs that control management services on the management node.

Additional integration utilities and tools

While you typically manage your storage with NetApp Element, NetApp Element API, and NetApp Element Plug-in for vCenter Server, you can use additional utilities and tools.

- [*NetApp Downloads: SolidFire vRO*](#)

Provides a convenient way to use the SolidFire API to administer your SolidFire storage system with VMware vRealize Orchestrator™.

- [*NetApp Downloads: Element SDK*](#)

Enables you to manage your SolidFire cluster using these tools:

- SolidFire command line
- SolidFire Postman API testing suite: Enables programmers to use a collection of Postman functions that test SolidFire API calls.
- SolidFire PowerShell: Enables programmers to use a collection of Microsoft® Windows® PowerShell® functions that use the SolidFire API to manage a SolidFire storage system.
- SolidFire SDK Java: Enables programmers to integrate the SolidFire API with the Java™ programming language.
- SolidFire SDK .NET: Enables programmers to integrate the SolidFire API with the .NET programming platform.
- SolidFire SDK Python: Enables programmers to integrate the SolidFire API with the Python™ programming language.

- [*NetApp Downloads: SolidFire Storage Replication Adapter*](#)

Integrates with the VMware® Site Recovery Manager™ (SRM) to enable communication with replicated SolidFire storage clusters (arrays) and execute supported workflows.

- [*NetApp Downloads: SolidFire VSS Provider*](#)

Integrates VSS shadow copies with SolidFire snapshots and clones.

Networking

The network setup for a SolidFire system consists of switch and port requirements. The implementation of these depends on your system.

Related concepts

[*Switch configuration for clusters running Element software*](#) on page 13

The NetApp Element software system has certain switch requirements and best practices for optimal storage performance.

Related reference

[*Network port requirements*](#) on page 13

You might need to allow the following TCP ports through your datacenter's edge firewall so that you can manage the system remotely and allow clients outside of your datacenter to connect to resources. Some of these ports might not be required, depending on how you use the system.

Switch configuration for clusters running Element software

The NetApp Element software system has certain switch requirements and best practices for optimal storage performance.

Storage nodes require 10 or 25GbE Ethernet switches, depending on specific node hardware, for iSCSI storage services and node intra-cluster services communication. 1GbE switches can be used for these types of traffic:

- Management of the cluster and the nodes
- Intra-cluster management traffic between the nodes
- Traffic between the cluster nodes and the management node virtual machine

Best Practice: You should implement the following best practices when configuring Ethernet switches for cluster traffic:

- For non-storage traffic in the cluster, deploy a pair of 1GbE switches to provide high availability and load sharing.
- On the storage network switches, deploy switches in pairs and configure and utilize jumbo frames (an MTU size of 9216 bytes). This ensures a successful installation and eliminates storage network errors due to fragmented packets.

Network port requirements

You might need to allow the following TCP ports through your datacenter's edge firewall so that you can manage the system remotely and allow clients outside of your datacenter to connect to resources. Some of these ports might not be required, depending on how you use the system.

See the port requirements [here](#).

System setup

Before you can use your SolidFire storage system, you must install and configure the management node, configure the individual nodes, create a cluster, and add drives to the cluster.

The SolidFire storage system provisions storage using volumes. Volumes are block devices accessed over the network by iSCSI or Fibre Channel clients. Accounts enable clients to connect to volumes on a node. You must create accounts to be able to access the volumes on a node.

Your hardware must be racked, cabled, and powered on before you perform the system setup tasks. Instructions for setting up the hardware are included in the hardware shipment.

When setting up the SolidFire storage system, you must follow a specific order of operations to ensure that your nodes and clusters are configured correctly.

You can optionally set up Fibre Channel nodes, depending on your environment.

Related tasks

[Using basic options in the Element software UI](#) on page 27

The NetApp Element software web user interface (Element UI) enables you to monitor and perform common tasks on your SolidFire system.

[Working with user accounts](#) on page 65

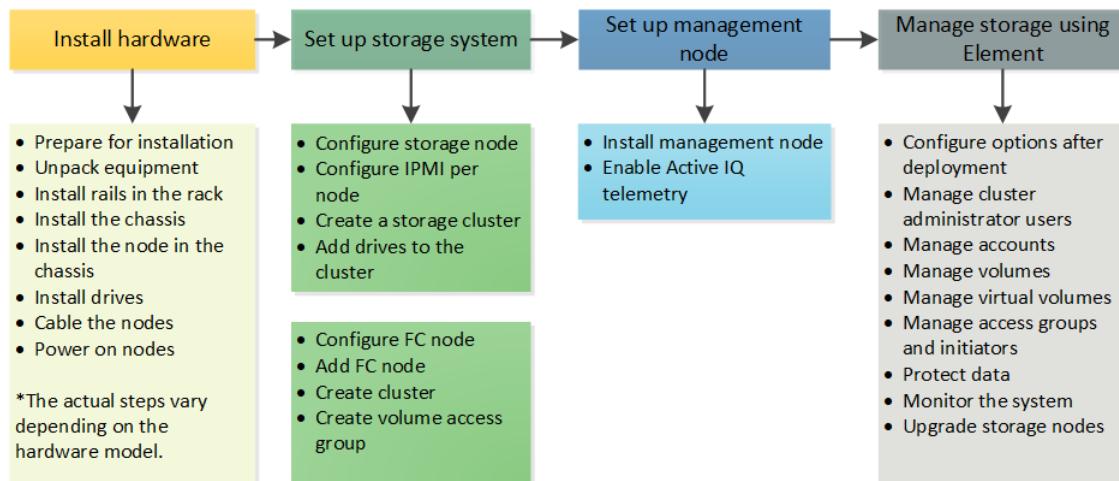
In SolidFire storage systems, user accounts enable clients to connect to volumes on a node. When you create a volume, it is assigned to a specific user account.

[Working with volumes](#) on page 68

The SolidFire system provisions storage using volumes. Volumes are block devices accessed over the network by iSCSI or Fibre Channel clients. From the Volumes page on the Management tab, you can create, modify, clone, and delete volumes on a node. You can also view statistics about volume bandwidth and I/O usage.

Setup overview

Before you get started, you might want to understand the sequence of installing and setting up NetApp Element software.



Determining which SolidFire components to install

You might want to check which SolidFire components, such as the management node, Active IQ and the NetApp Monitoring Agent (NMA), that you should install, depending on configuration and deployment choices.

About this task

The following table lists the additional components and indicates whether you should install them.

Component	Standalone SolidFire storage cluster	NetApp HCI cluster
Management node	Recommended	Installed by default, required
Active IQ	Recommended*	Recommended*
NetApp Monitoring Agent	Not supported	Recommended

*Active IQ is required for capacity-licensed SolidFire storage clusters.

Steps

1. Determine which components should be installed.
2. Complete the installation according to the [install the management node](#) procedure.

Note: To set up Active IQ, use the `--telemetry_active` parameter in the setup script to enable data collection for analytics by Active IQ.

3. For NetApp Monitoring Agent information, see the deployment information.

[NetApp HCI Documentation Center](#)

Setting up an Element storage system

Setting up a NetApp Element software storage system involves configuring a storage node, creating a storage cluster, and adding drives to the cluster. If you use a Fibre Channel network, you can configure a Fibre Channel node.

Steps

1. [Configuring a storage node](#) on page 16

You must configure individual nodes before you can add them to a cluster. After you install and cable a node in a rack unit and power it on, you can configure the node network settings using the per-node UI or the node terminal user interface (TUI). Ensure that you have the necessary network configuration information for the node before proceeding.

2. [Creating a storage cluster](#) on page 19

You can create a storage cluster after you have configured all of the individual nodes. When you create a cluster, a cluster administrator user account is automatically created for you. The cluster administrator has permission to manage all cluster attributes and can create other cluster administrator accounts.

3. [Accessing the Element software user interface](#) on page 20

You can access the Element UI by using the management virtual IP (MVIP) address of the primary cluster node.

4. [Adding drives to a cluster](#) on page 20

When you add a node to the cluster or install new drives in an existing node, the drives automatically register as available. You must add the drives to the cluster by using either the Element UI or API before they can participate in the cluster.

5. [Configuring a Fibre Channel node](#) on page 21

Fibre Channel nodes enable you to connect the cluster to a Fibre Channel network fabric. Fibre Channel nodes are added in pairs, and operate in active-active mode (all nodes actively process traffic for the cluster). Clusters running Element software version 9.0 and later support up to four nodes; clusters running previous versions support a maximum of two nodes.

6. *Setting up a management node* on page 23

You can install the NetApp Element software management node (mNode) to upgrade and provide system services, manage cluster assets and settings, run system tests and utilities, and enable NetApp Support access for troubleshooting.

Configuring a storage node

You must configure individual nodes before you can add them to a cluster. After you install and cable a node in a rack unit and power it on, you can configure the node network settings using the per-node UI or the node terminal user interface (TUI). Ensure that you have the necessary network configuration information for the node before proceeding.

There are two options for configuring storage nodes:

Per-node UI

Use the per-node UI (`https://<node management IP>:442`) to configure node network settings.

Note: Use the DHCP 1G management IP address displayed in the menu bar at the top of the TUI to access the per-node UI.

TUI

Use the node terminal user interface (TUI) to configure the node.

You cannot add a node with DHCP-assigned IP addresses to a cluster. You can use the DHCP IP address to initially configure the node in the per-node UI, TUI, or API. During this initial configuration, you can add static IP address information so that you can add the node to a cluster.

After initial configuration, you can access the node using the node's management IP address. You can then change the node settings, add it to a cluster, or use the node to create a cluster. You can also configure a new node using Element software API methods.

Note: Beginning in Element version 11.0, nodes can be configured with IPv4, IPv6, or both addresses for their management network. This applies to both storage nodes and management nodes, except for management node 11.3 and later which does not support IPv6. When you create a cluster, only a single IPv4 or IPv6 address can be used for the MVIP and the corresponding address type must be configured on all nodes.

Related tasks

[*Configuring the storage node using the per-node UI* on page 17](#)

You can configure nodes using the per-node user interface.

[*Configuring the node using the TUI* on page 18](#)

You can use the terminal user interface (TUI) to perform initial configuration for new nodes.

[*Creating a storage cluster* on page 19](#)

You can create a storage cluster after you have configured all of the individual nodes. When you create a cluster, a cluster administrator user account is automatically created for you. The cluster administrator has permission to manage all cluster attributes and can create other cluster administrator accounts.

Related reference

[*Node states* on page 18](#)

A node can be in one of several states depending on the level of configuration.

Related information

[NetApp SolidFire Installation](#)

Configuring the storage node using the per-node UI

You can configure nodes using the per-node user interface.

About this task

- You can configure the node to have either an IPv4 or IPv6 address.
- You need the DHCP address displayed in the TUI to access a node. You cannot use DHCP addresses to add a node to a cluster.



Attention: You should configure the management (Bond1G) and storage (Bond10G) interfaces for separate subnets. Bond1G and Bond10G interfaces configured for the same subnet cause routing problems when storage traffic is sent via the Bond1G interface. If you must use the same subnet for management and storage traffic, manually configure management traffic to use the Bond10G interface. You can do this for each node using the **Cluster Settings** page of the per-node UI.

Steps

1. In a browser window, enter the DHCP IP address of a node.
You must add the extension :442 to access the node; for example, `https://172.25.103.6:442`.

The **Network Settings** tab opens with the **Bond1G** section.

2. Enter the 1G management network settings.
3. Click **Apply Changes**.
4. Click **Bond10G** to display the 10G storage network settings.
5. Enter the 10G storage network settings.
6. Click **Apply Changes**.
7. Click **Cluster Settings**.
8. Enter the hostname for the 10G network.
9. Enter the cluster name.

Important: This name must be added to the configuration for all nodes before a cluster can be created. All the nodes in a cluster must have identical cluster names. Cluster names are case-sensitive.

10. Click **Apply Changes**.

Related tasks

[Configuring the node using the TUI](#) on page 18

You can use the terminal user interface (TUI) to perform initial configuration for new nodes.

Related reference

[Network settings details from the per-node UI](#) on page 164

You can change the storage node network settings to give the node a new set of network attributes.

[Cluster settings details from the per-node UI](#) on page 165

You can verify cluster settings for a storage node after cluster configuration and modify the node hostname.

Configuring the node using the TUI

You can use the terminal user interface (TUI) to perform initial configuration for new nodes.

About this task

You should configure the Bond1G (Management) and Bond10G (Storage) interfaces for separate subnets. Bond1G and Bond10G interfaces configured for the same subnet causes routing problems when storage traffic is sent via the Bond1G interface. If you must use the same subnet for management and storage traffic, manually configure management traffic to use the Bond10G interface. You can do this for each node using the **Cluster > Nodes** page of the Element UI.

Steps

1. Attach a keyboard and monitor to the node and then power on the node.

The NetApp Storage Main menu of the TUI appears on the tty1 terminal.

Note: If the node cannot reach your configuration server, the TUI displays an error message.

Check your configuration server connection or the networking connection to resolve the error.

2. Select **Network > Network Config.**

Tip: To navigate through the menu, press the Up or Down arrow keys. To move to another button or to the fields from the buttons, press **Tab**. To navigate between fields, use the Up or Down arrow keys.

3. Select **Bond1G (Management)** or **Bond10G (Storage)** to configure the 1G and 10G network settings for the node.
4. For the Bond mode and Status fields, press **Tab** to select the Help button and identify the available options.

All the nodes in a cluster must have identical cluster names. Cluster names are case-sensitive. If a DHCP server is running on the network with available IP addresses, the 1GbE address appears in the Address field.

5. Press **Tab** to select the **OK** button and save the changes.

The node is put in a pending state and can be added to an existing cluster or a new cluster.

Related tasks

[Configuring the storage node using the per-node UI](#) on page 17

You can configure nodes using the per-node user interface.

Node states

A node can be in one of several states depending on the level of configuration.

Available

The node has no associated cluster name and is not yet part of a cluster.

Pending

The node is configured and can be added to a designated cluster.

Authentication is not required to access the node.

Pending Active

The system is in the process of installing compatible Element software on the node. When complete, the node will move to the **Active** state.

Active

- The node is participating in a cluster.
- Authentication is required to modify the node.

In each of these states, some fields are read only.

Creating a storage cluster

You can create a storage cluster after you have configured all of the individual nodes. When you create a cluster, a cluster administrator user account is automatically created for you. The cluster administrator has permission to manage all cluster attributes and can create other cluster administrator accounts.

Before you begin

- You have installed the management node.
- You have configured all of the individual nodes.

About this task

During new node configuration, 1G or 10G Management IP (MIP) addresses are assigned to each node. You must use one of the node IP addresses created during configuration to open the Create a New Cluster page. The IP address you use depends on the network you have chosen for cluster management.

Note: When creating a new cluster:

- If you are using storage nodes that reside in a shared chassis, you might want to consider designing for chassis-level failure protection using the protection domains feature.
- If a shared chassis is not in use, you can define a custom protection domain layout.

Steps

- In a browser window, enter a node MIP address.
- In **Create a New Cluster**, enter the following information:
 - Management VIP: Routable virtual IP on the 1GbE or 10GbE network for network management tasks.
Note: You can create a new cluster using IPv4 or IPv6 addressing.
 - iSCSI (storage) VIP: Virtual IP on the 10GbE network for storage and iSCSI discovery.
Note: You cannot change the MVIP, SVIP, or cluster name after you create the cluster.
 - User name: The primary cluster administrator user name for authenticated access to the cluster. You must save the user name for future reference.
Note: You can use uppercase and lowercase letters, special characters, and numbers for the user name and password.
 - Password: Password for authenticated access to the cluster. You must save the password for future reference.
- Two-way data protection is enabled by default. You cannot change this setting.
- Read the End User License Agreement, and click **I Agree**.
- Optional: In the Nodes list, ensure that the check boxes for nodes that should not be included in the cluster are not selected.
- Click **Create Cluster**.

The system might take several minutes to create the cluster depending on the number of nodes in the cluster. On a properly configured network, a small cluster of five nodes should take less than one minute. After the cluster is created, the Create a New Cluster window is redirected to the MVIP URL address for the cluster and displays the Element UI.

Related information

[Managing storage with the Element API](#)

Accessing the Element software user interface

You can access the Element UI by using the management virtual IP (MVIP) address of the primary cluster node.

Before you begin

You must ensure that popup blockers and NoScript settings are disabled in your browser.

About this task

You can access the UI using IPv4 or IPv6 addressing, depending on configuration during cluster creation.

Steps

1. Choose one of the following:

- IPv6: Enter `https://[IPv6 MVIP address]` For example:

`https://[fd20:8b1e:b256:45a::1234]/`

- IPv4: Enter `https://<IPv4 MVIP address>` For example:

`https://10.123.456.789/`

2. For DNS, enter the host name.

3. Click through any authentication certificate messages.

Adding drives to a cluster

When you add a node to the cluster or install new drives in an existing node, the drives automatically register as available. You must add the drives to the cluster by using either the Element UI or API before they can participate in the cluster.

About this task

Drives are not displayed in the Available Drives list when the following conditions exist:

- Drives are in Active, Removing, Erasing, or Failed state.
- The node of which the drive is a part of is in Pending state.

Steps

1. From the Element user interface, select **Cluster > Drives**.

2. Click **Available** to view the list of available drives.

3. Do one of the following:

- To add individual drives, click the **Actions** icon for the drive you want to add and click **Add**.
- To add multiple drives, select the check boxes of the drives to add, click **Bulk Actions**, and click **Add**.

Related information

[How to calculate max provisioned space in a SolidFire cluster](#)

Configuring a Fibre Channel node

Fibre Channel nodes enable you to connect the cluster to a Fibre Channel network fabric. Fibre Channel nodes are added in pairs, and operate in active-active mode (all nodes actively process traffic for the cluster). Clusters running Element software version 9.0 and later support up to four nodes; clusters running previous versions support a maximum of two nodes.

You must ensure that the following conditions are met before you configure a Fibre Channel node:

- At least two Fibre Channel nodes are connected to Fibre Channel switches.
- All SolidFire Fibre Channel ports should be connected to your Fibre Channel fabric. The four SolidFire Bond10G network connections should be connected in one LACP bond group at the switch level. This will enable the best overall performance from the Fibre Channel systems.
- Review and validate all best practices for Fibre Channel clusters included in this NetApp Knowledge Base article.

SolidFire FC cluster best practice

Network and cluster configuration steps are the same for Fibre Channel nodes and storage nodes.

When you create a new cluster with Fibre Channel nodes and SolidFire storage nodes, the worldwide port name (WWPN) addresses for the nodes are available in the Element UI. You can use the WWPN addresses to zone the Fibre Channel switch.

WWPNs are registered in the system when you create a new cluster with nodes. In the Element UI, you can find the WWPN addresses from the WWPN column of the FC Ports tab, which you access from the Cluster tab.

Related tasks

Configuring the storage node using the per-node UI on page 17

You can configure nodes using the per-node user interface.

Configuring the node using the TUI on page 18

You can use the terminal user interface (TUI) to perform initial configuration for new nodes.

Creating a storage cluster on page 19

You can create a storage cluster after you have configured all of the individual nodes. When you create a cluster, a cluster administrator user account is automatically created for you. The cluster administrator has permission to manage all cluster attributes and can create other cluster administrator accounts.

Related information

SolidFire Fibre Channel Configuration Guide

Adding Fibre Channel nodes to a cluster

You can add Fibre Channel nodes to a cluster when more storage is needed or during cluster creation. Fibre Channel nodes require initial configuration when they are first powered on. After the node is configured, it appears in the list of pending nodes and you can add it to a cluster.

About this task

The software version on each Fibre Channel node in a cluster must be compatible. When you add a Fibre Channel node to a cluster, the cluster installs the cluster version of Element on the new node as needed.

Steps

1. Select **Cluster > Nodes**.
2. Click **Pending** to view the list of pending nodes.
3. Do one of the following:
 - To add individual nodes, click the **Actions** icon for the node you want to add.

- To add multiple nodes, select the check box of the nodes to add, and then **Bulk Actions**.

Note:

If the node you are adding has a different version of Element than the version running on the cluster, the cluster asynchronously updates the node to the version of Element running on the cluster master. After the node is updated, it automatically adds itself to the cluster. During this asynchronous process, the node will be in a pendingActive state.

4. Click **Add.**

The node appears in the list of active nodes.

Creating a new cluster with Fibre Channel nodes

You can create a new cluster after you have configured the individual Fibre Channel nodes. When you create a cluster, a cluster administrator user account is automatically created for you. The cluster administrator has permission to manage all cluster attributes and can create other cluster administrator accounts.

Before you begin

You have configured the individual Fibre Channel nodes.

About this task

During new node configuration, 1G or 10G Management IP (MIP) addresses are assigned to each node. You must use one of the node IP addresses created during configuration to open the Create a New Cluster page. The IP address you use depends on the network you have chosen for cluster management.

Steps

1. In a browser window, enter a node MIP address.
2. In **Create a New Cluster**, enter the following information:
 - Management VIP: Routable virtual IP on the 1GbE or 10GbE network for network management tasks.
 - iSCSI (storage) VIP: Virtual IP on the 10GbE network for storage and iSCSI discovery.

Note: You cannot change the SVIP after you create the cluster.

 - User name: The primary Cluster Admin user name for authenticated access to the cluster. You must save the user name for future reference.

Note: You can use uppercase and lowercase letters, special characters, and numbers for the user name.

 - Password: Password for authenticated access to the cluster. You must save the user name for future reference.

Two-way data protection is enabled by default. You cannot change this setting.

3. Read the End User License Agreement, and click **I Agree**.
4. Optional: In the Nodes list, ensure that the check boxes for nodes that should not be included in the cluster are not selected.
5. Click **Create Cluster**.

The system might take several minutes to create the cluster depending on the number of nodes in the cluster. On a properly configured network, a small cluster of five nodes should take less than one minute. After the cluster is created, the Create a New Cluster window is redirected to the MVIP URL address for the cluster and displays the web UI.

Zoning for Fibre Channel nodes

When you create a new cluster with Fibre Channel nodes and SolidFire storage nodes, the worldwide port name (WWPN) addresses for the nodes are available in the web UI. You can use the WWPN addresses to zone the Fibre Channel switch.

WWPNs are registered in the system when you create a new cluster with nodes. In the Element UI, you can find the WWPN addresses from the WWPN column of the FC Ports tab, which you access from the Cluster tab.

Creating a volume access group for Fibre Channel clients

Volume access groups enable communication between Fibre Channel clients and volumes on a SolidFire storage system. Mapping Fibre Channel client initiators (WWPN) to the volumes in a volume access group enables secure data I/O between a Fibre Channel network and a SolidFire volume.

About this task

You can also add iSCSI initiators to a volume access group; this gives the initiators access to the same volumes in the volume access group.

Steps

1. Click **Management > Access Groups**.
2. Click **Create Access Group**.
3. Enter a name for the volume access group in the **Name** field.
4. Select and add the Fibre Channel initiators from the **Unbound Fibre Channel Initiators** list.
Note: You can add or delete initiators at a later time.
5. Optional: Select and add an iSCSI initiator from the **Initiators** list.
6. To attach volumes to the access group, perform the following steps:
 - a. Select a volume from the **Volumes** list.
 - b. Click **Attach Volume**.
7. Click **Create Access Group**.

Setting up a management node

You can install the NetApp Element software management node (mNode) to upgrade and provide system services, manage cluster assets and settings, run system tests and utilities, and enable NetApp Support access for troubleshooting.

Step

See the [install the management node](#) documentation.

Note: To set up Active IQ, use the `--telemetry_active` parameter in the setup script to enable data collection for analytics by Active IQ.

Configuring SolidFire system options after deployment

After you set up your SolidFire system, you might want to perform some optional tasks. Additionally, you can configure settings for multi-factor authentication, external key management, and Federal Information Processing Standards (FIPS) security.

Related concepts

[Enabling Multi-factor authentication](#) on page 41

Multi-factor authentication (MFA) uses a third-party Identity Provider (IdP) via the Security Assertion Markup Language (SAML) to manage user sessions. MFA enables administrators to configure additional factors of authentication as required, such as password and text message, and password and email message.

[Getting started with External key management](#) on page 62

External key management (EKM) provides secure Authentication Key (AK) management in conjunction with an off-cluster external key server (EKS). The EKS provides secure generation and storage of the AKs.

Related tasks

[Creating a cluster supporting FIPS drives](#) on page 57

Security is becoming increasingly critical for the deployment of solutions in many customer environments. Federal Information Processing Standards (FIPS) are standards for computer security and interoperability. FIPS 140-2 certified encryption for data at rest is a component of the overall security solution.

Changing the Element software default SSL certificate

You can change the default SSL certificate and private key of the storage node in the cluster using the NetApp Element API.

When a NetApp Element software cluster is created, the cluster creates a unique self-signed Secure Sockets Layer (SSL) certificate and private key that is used for all HTTPS communication via the Element UI, per-node UI, or APIs. Element software supports self-signed certificates as well as certificates that are issued and verified by a trusted Certificate Authority (CA).

You can use the following API methods to get more information about the default SSL certificate and make changes. For information about each method, see the *NetApp Element Software API Reference Guide*.

GetSSLCertificate

You can use this method to retrieve information about the currently installed SSL certificate including all certificate details.

SetSSLCertificate

You can use this method to set the cluster and per-node SSL certificates to the certificate and private key you supply. The system validates the certificate and private key to prevent an invalid certificate from being applied.

RemoveSSLCertificate

This method removes the currently installed SSL certificate and private key. The cluster then generates a new self-signed certificate and private key.

Note: The cluster SSL certificate is automatically applied to all new nodes added to the cluster. Any node removed from the cluster reverts to a self-signed certificate and all user-defined certificate and key information is removed from the node.

Upgrading storage nodes

You can upgrade the Element software on the storage nodes of your cluster using the HealthTools suite. You can use HealthTools from a connected or dark site. You must use management node 11.0, 11.1, or later to use the latest HealthTools.

To update your Element storage as part of an end-to-end system upgrade, follow the upgrade sequence for your system: [NetApp Documentation: Upgrades overview](#)

Important: If you are upgrading an H610S series node to Element 12.0 or later, there are additional steps required to complete the storage upgrade for each node. See the following topic for H610S procedures: [NetApp Documentation: Upgrade Element Storage](#)

Updating management services for Element storage-based installations

As of the management node 11.3 release, the management node functions as a microservice host, allowing for quicker updates of select software services outside of major releases for NetApp HCI and SolidFire all-flash storage. These microservices or management services are updated frequently as service bundles from an online software repository. You can use Hybrid Cloud Control to update management services to the latest version. Alternatively, you can use management services REST APIs that are available from the management node to stay up-to-date.

To update your management services as part of an end-to-end Element storage system upgrade:

[*NetApp Documentation: Upgrades overview*](#)

To update your management services outside of a major Element storage release: [*NetApp*](#)

[*Documentation: Update management services*](#)

Using basic options in the Element software UI

The NetApp Element software web user interface (Element UI) enables you to monitor and perform common tasks on your SolidFire system.

About this task

You can search for information with filters, sort lists, see API commands activated by UI activity, and provide feedback.

Related tasks

[Limiting results by using filters](#) on page 28

You can filter list information on pages in the Element UI. When viewing lists (such as volumes and snapshots), you can add one or more filters to focus the information and make it more easily fit on the page.

[Viewing API activity](#) on page 28

The Element system uses the NetApp Element API as the foundation for its features and functionality. The Element UI enables you to view various types of real-time API activity on the system as you use the interface. With the API log, you can view user-initiated and background system API activity, as well as API calls made on the page you are currently viewing.

[Sorting lists](#) on page 28

You can sort list information by one or more criteria on certain pages in the Element UI. This helps you arrange the information you need on the screen.

[Providing feedback](#) on page 30

You can help improve the Element software web user interface and address any UI issues by using the feedback form that is accessible throughout the UI.

Related reference

[Icons in the Element interface](#) on page 29

The NetApp Element software interface displays icons to represent actions you can take on system resources.

Accessing the Element software user interface

You can access the Element UI by using the management virtual IP (MVIP) address of the primary cluster node.

Before you begin

You must ensure that popup blockers and NoScript settings are disabled in your browser.

About this task

You can access the UI using IPv4 or IPv6 addressing, depending on configuration during cluster creation.

Steps

1. Choose one of the following:

- IPv6: Enter `https://[IPv6 MVIP address]` For example:

`https://[fd20:8b1e:b256:45a::1234]/`

- IPv4: Enter `https://<IPv4 MVIP address>` For example:

`https://10.123.456.789/`

2. For DNS, enter the host name.

3. Click through any authentication certificate messages.

Limiting results by using filters

You can filter list information on pages in the Element UI. When viewing lists (such as volumes and snapshots), you can add one or more filters to focus the information and make it more easily fit on the page.

Steps

1. When viewing list information, click **Filter**.
2. Expand the **Filter By** field.
3. Choose a column to filter by from the leftmost element in the field.
4. Choose a constraint for the column.
5. Enter text to filter by.
6. Click **Add**.
The system runs the new filter on the information in the list, and temporarily stores the new filter in the Filter By field.
7. Optional: To add another filter, click **Add** and again choose the filter.
8. Optional: Click **Clear All** to remove the list of filters and display the unfiltered list information.

Sorting lists

You can sort list information by one or more criteria on certain pages in the Element UI. This helps you arrange the information you need on the screen.

Steps

1. To sort on a single column, click the column heading until the information is sorted in the desired order.
2. To sort using multiple columns, click the column heading for each column you want to sort by until the information in each column is sorted in the desired order.
The Sort button appears when you sort using multiple columns.
3. To reorder the sort criteria, perform the following steps:
 - a. Click **Sort**.
The system populates the Sort By field with your column selections.
 - b. Arrange the columns in the Sort By field in the order you want the list to be sorted.
The system sorts the list information.
4. To remove a single sort criterion, click the **Remove** icon next to the name of a sort criteria.
5. Optional: To remove all sort criteria, click **Clear All**.

Viewing API activity

The Element system uses the NetApp Element API as the foundation for its features and functionality. The Element UI enables you to view various types of real-time API activity on the system as you use the interface. With the API log, you can view user-initiated and background system API activity, as well as API calls made on the page you are currently viewing.

About this task

You can use the API log to identify what API methods are used for certain tasks, and see how to use the API methods and objects to build custom applications. For information about each method, see the *NetApp Element Software API Reference Guide*.

Steps

1. From the Element UI navigation bar, click **API Log**.
2. To modify the type of API activity displayed in the **API Log** window, perform the following steps:
 - a. Select **Requests** to display API request traffic.
 - b. Select **Responses** to display API response traffic.
 - c. Filter the types of API traffic by selecting one of the following:
 - **User Initiated**: API traffic by your activities during this web UI session.
 - **Background Polling**: API traffic generated by background system activity.
 - **Current Page**: API traffic generated by tasks on the page you are currently viewing.

Related information

[Managing storage with the Element API](#)

Interface refresh rate impacted by cluster load

Depending on API response times, the cluster might automatically adjust the data refresh interval for certain portions of the NetApp Element software page you are viewing.

The refresh interval is reset to the default when you reload the page in your browser. You can see the current refresh interval by clicking the cluster name in the upper-right of the page. Note that the interval controls how often API requests are made, not how quickly the data comes back from the server.

When a cluster is under heavy load, it might queue API requests from the Element UI. In rare circumstances, when system response is significantly delayed, such as a slow network connection combined with a busy cluster, you might be logged out of the Element UI if the system does not respond to queued API requests quickly enough. If you are redirected to the logout screen, you can log in again after dismissing any initial browser authentication prompt. Upon returning to the overview page, you might be prompted for cluster credentials if they are not saved by your browser.

Icons in the Element interface

The NetApp Element software interface displays icons to represent actions you can take on system resources.

The following table provides a quick reference:

Icon	Description
	Actions
	Backup to
	Clone or copy
	Delete or purge
	Edit

Icon	Description
	Filter
	Pair
	Refresh
	Restore
	Restore from
	Rollback
	Snapshot

Providing feedback

You can help improve the Element software web user interface and address any UI issues by using the feedback form that is accessible throughout the UI.

Steps

1. From any page in the Element UI, click the **Feedback** button.
2. Enter relevant information in the Summary and Description fields.
3. Attach any helpful screenshots.
4. Enter a name and email address.
5. Select the check box to include data about your current environment.
6. Click **Submit**.

System management

You can manage your system in the Element UI. This includes creating and managing cluster administrators, managing cluster settings, and upgrading software.

Related concepts

[Managing cluster administrator user accounts](#) on page 31

You can manage cluster administrator accounts for a SolidFire storage system. Available cluster administrator management functions include creating, deleting, and editing cluster administrator accounts, changing the cluster administrator password, and configuring LDAP settings to manage system access for users.

[Configuring cluster settings](#) on page 43

You can view and change cluster-wide settings and perform cluster-specific tasks from the Cluster tab of the Element UI.

Managing cluster administrator user accounts

You can manage cluster administrator accounts for a SolidFire storage system. Available cluster administrator management functions include creating, deleting, and editing cluster administrator accounts, changing the cluster administrator password, and configuring LDAP settings to manage system access for users.

Related concepts

[Storage cluster administrator account types](#) on page 32

There are two types of administrator accounts that can exist in a storage cluster running NetApp Element software: the primary cluster administrator account and a cluster administrator account.

Related tasks

[Creating a cluster administrator account](#) on page 33

You can create new cluster administrator accounts with permissions to allow or restrict access to specific areas of the storage system. When you set cluster administrator account permissions, the system grants read-only rights for any permissions you do not assign to the cluster administrator.

[Editing cluster administrator permissions](#) on page 34

You can change cluster administrator account privileges for reporting, nodes, drives, volumes, accounts, and cluster-level access. When you enable a permission, the system assigns write access for that level. The system grants the administrator user read-only access for the levels that you do not select.

[Changing passwords for cluster administrator accounts](#) on page 34

You can use the Element UI to change cluster administrator passwords.

[Enabling LDAP authentication with the Element user interface](#) on page 35

You can configure storage system integration with an existing LDAP server. This enables LDAP administrators to centrally manage storage system access for users.

[Disabling LDAP](#) on page 41

You can disable LDAP integration using the Element UI.

Related reference

[Cluster Admins details](#) on page 32

On the Cluster Admins page of the Users tab, you can view the following information.

Storage cluster administrator account types

There are two types of administrator accounts that can exist in a storage cluster running NetApp Element software: the primary cluster administrator account and a cluster administrator account.

Primary cluster administrator account

This administrator account is created when the cluster is created. This account is the primary administrative account with the highest level of access to the cluster. This account is analogous to a root user in a Linux system. You can change the password for this administrator account.

Cluster administrator account

You can give a cluster administrator account a limited range of administrative access to perform specific tasks within a cluster. The credentials assigned to each cluster administrator account are used to authenticate API and Element UI requests within the storage system.

Note: A local (non-LDAP) cluster administrator account is required to access active nodes in a cluster via the per-node UI. Account credentials are not required to access a node that is not yet part of a cluster.

Cluster Admins details

On the Cluster Admins page of the Users tab, you can view the following information.

ID

Sequential number assigned to the cluster administrator account.

Username

The name given to the cluster administrator account when it was created.

Access

The user permissions assigned to the user account. Possible values:

- read
- reporting
- nodes
- drives
- volumes
- accounts
- clusterAdmins
- administrator

Note: All permissions are available to the administrator access type.

Type

The type of cluster administrator. Possible values:

- Cluster
- Ldap

Attributes

If the cluster administrator account was created using the Element API, this column shows any name-value pairs that were set using that method. See the *NetApp Element Software API Reference Guide*.

Related information

[Managing storage with the Element API](#)

Creating a cluster administrator account

You can create new cluster administrator accounts with permissions to allow or restrict access to specific areas of the storage system. When you set cluster administrator account permissions, the system grants read-only rights for any permissions you do not assign to the cluster administrator.

Before you begin

If you want to create an LDAP cluster administrator account, ensure that LDAP is configured on the cluster before you begin.

About this task

You can later change cluster administrator account privileges for reporting, nodes, drives, volumes, accounts, and cluster-level access. When you enable a permission, the system assigns write access for that level. The system grants the administrator user read-only access for the levels that you do not select.

You can also later remove any cluster administrator user account created by a system administrator. You cannot remove the primary cluster administrator account that was created when the cluster was created.

Steps

1. To create a cluster-wide (non-LDAP) cluster administrator account, perform the following actions:
 - a. Click **Users > Cluster Admins**.
 - b. Click **Create Cluster Admin**.
 - c. Select the **Cluster** user type.
 - d. Enter a user name and password for the account and confirm password.
 - e. Select user permissions to apply to the account.
 - f. Select the check box to agree to the End User License Agreement.
 - g. Click **Create Cluster Admin**.
2. To create a cluster administrator account in the LDAP directory, perform the following actions:
 - a. Click **Cluster > LDAP**.
 - b. Ensure that LDAP Authentication is enabled.
 - c. Click **Test User Authentication** and copy the distinguished name that appears for the user or one of the groups of which the user is a member so that you can paste it later.
 - d. Click **Users > Cluster Admins**.
 - e. Click **Create Cluster Admin**.
 - f. Select the **LDAP** user type.
 - g. In the Distinguished Name field, follow the example in the text box to enter a full distinguished name for the user or group. Alternatively, paste it from the distinguished name you copied earlier.

If the distinguished name is part of a group, then any user that is a member of that group on the LDAP server will have permissions of this admin account.

To add LDAP Cluster Admin users or groups the general format of the username is “LDAP:<Full Distinguished Name>”.

- h. Select user permissions to apply to the account.
- i. Select the check box to agree to the End User License Agreement.
- j. Click **Create Cluster Admin**.

Related tasks

[Enabling LDAP authentication with the Element user interface](#) on page 35

You can configure storage system integration with an existing LDAP server. This enables LDAP administrators to centrally manage storage system access for users.

Editing cluster administrator permissions

You can change cluster administrator account privileges for reporting, nodes, drives, volumes, accounts, and cluster-level access. When you enable a permission, the system assigns write access for that level. The system grants the administrator user read-only access for the levels that you do not select.

Steps

1. Click **Users > Cluster Admins**.
2. Click the Actions icon for the cluster administrator you want to edit.
3. Click **Edit**.
4. Select user permissions to apply to the account.
5. Click **Save Changes**.

Changing passwords for cluster administrator accounts

You can use the Element UI to change cluster administrator passwords.

About this task

Note: To change the password for the cluster administrator in the management node, see management node documentation.

[Management node overview](#)

Steps

1. Click **Users > Cluster Admins**.
2. Click the Actions icon for the cluster administrator you want to edit.
3. Click **Edit**.
4. In the Change Password field, enter a new password and confirm it.
5. Click **Save Changes**.

Related concepts

[Storage cluster administrator account types](#) on page 32

There are two types of administrator accounts that can exist in a storage cluster running NetApp Element software: the primary cluster administrator account and a cluster administrator account.

Managing LDAP

You can set up the Lightweight Directory Access Protocol (LDAP) to enable secure, directory-based login functionality to SolidFire storage. You can configure LDAP at the cluster level and authorize LDAP users and groups.

Managing LDAP involves setting up LDAP authentication to a SolidFire cluster using an existing Microsoft Active Directory environment and testing the configuration.

Note: You can use both IPv4 and IPv6 addresses.

Enabling LDAP involves the following high-level steps, each described in detail:

1. **Verify inputs.** Validate that you have all of the details required to configure LDAP authentication.

2. **Enable LDAP authentication.** Use either the Element UI or the Element API.
3. **Validate the LDAP configuration.** Optionally, check that the cluster is configured with the correct values by running the `GetLdapConfiguration` API method or by checking the LCAP configuration using the Element UI.
4. **Test the LDAP authentication** (with the `readonly` user). Test that the LDAP configuration is correct either by running the `TestLdapAuthentication` API method or by using the Element UI. For this initial test, use the username “`SAMAccountName`” of the `readonly` user. This will validate that your cluster is configured correctly for LDAP authentication and also validate that the `readonly` credentials and access are correct. If this step fails, repeat steps 1 through 3.
5. **Test the LDAP authentication** (with a user account that you want to add). Repeat step 4 with a user account that you want to add as an Element cluster admin. Copy the distinguished name (DN) or the user (or the group). This DN will be used in step 6.
6. **Add the LDAP cluster admin** (copy and paste the DN from the Test LDAP authentication step). Using either the Element UI or the `AddLdapClusterAdmin` API method, create a new cluster admin user with the appropriate access level. For the username, paste in the full DN you copied in Step 5. This assures that the DN is formatted correctly.
7. **Test the cluster admin access.** Log in to the cluster using the newly created LDAP cluster admin user. If you added an LDAP group, you can log in as any user in that group.

Completing pre-configuration steps for LDAP support

Before you enable LDAP support in Element, you should set up a Windows Active Directory Server and perform other pre-configuration tasks.

Steps

1. Set up a Windows Active Directory Server.
2. Optional: Enable LDAPS support.
3. Create users and groups.
4. Create a read-only service account (such as "sfreadonly") to be used for searching the LDAP directory.

Enabling LDAP authentication with the Element user interface

You can configure storage system integration with an existing LDAP server. This enables LDAP administrators to centrally manage storage system access for users.

About this task

You can configure LDAP with either the Element user interface or the Element API. This procedure describes how to configure LDAP using the Element UI.

This example shows how to configure LDAP authentication on SolidFire and it uses `SearchAndBind` as the authentication type. The example uses a single Windows Server 2012 R2 Active Directory Server.

Steps

1. Click **Cluster > LDAP**.
2. Click **Yes** to enable LDAP authentication.
3. Click **Add a Server**.
4. Enter the **Host Name/IP Address**.

Note: An optional custom port number can also be entered.

For example, to add a custom port number, enter <host name or ip address>:<port number>

5. Optional: Select **Use LDAPS Protocol**.
6. Enter the required information in **General Settings**.

[LDAP details](#) on page 39

The screenshot shows two configuration pages for LDAP settings. The top section, 'LDAP Servers', contains a table with one row: 'Host Name/IP Address' (192.168.9.99), a 'Remove' button, and a checked 'Use LDAPS Protocol' checkbox. Below this is a green 'Add a Server' button. The bottom section, 'General Settings', contains several configuration fields: 'Auth Type' (Search and Bind), 'Search Bind DN' (mwhite@thewhites.ca), 'Search Bind Password' (e.g. password), a 'Show password' checkbox (unchecked), 'User Search Base DN' (OU=Home users,DC=thewhites,DC=ca), 'User Search Filter' (&(objectClass=person)(|(sAMAccountName=%USER)), 'Group Search Type' (Active Directory), and 'Group Search Base DN' (OU=Home users,DC=thewhites,DC=ca). At the bottom of this section is a 'Save Changes' button.

7. Click **Enable LDAP**.
8. Click **Test User Authentication** if you want to test the server access for a user.
9. Copy the distinguished name and user group information that appears for use later when creating cluster administrators.
10. Optional: Click **Save Changes** to save any new settings.
11. To create a user in this group so that anyone can log in, complete the following:
 - a. Click **User > View**.

Create a New Cluster Admin

Select User Type

Cluster LDAP

Enter User Details

Distinguished Name

CN=StorageAdmins,OU=Home
users,DC=thewhites,DC=ca

Select User Permissions

Reporting Volumes
 Nodes Accounts
 Drives Cluster Admin

Accept the Following End User License Agreement

- b. For the new user, click **LDAP** for the User Type, and paste the group you copied to the Distinguished Name field.
- c. Select the permissions, typically all permissions.
- d. Scroll down to the End User License Agreement and click **I accept**.
- e. Click **Create Cluster Admin**.

Now you have a user with the value of an Active Directory group.

After you finish

To test this, log out of the Element UI and log back in as a user in that group.

Related tasks

[Creating a cluster administrator account](#) on page 33

You can create new cluster administrator accounts with permissions to allow or restrict access to specific areas of the storage system. When you set cluster administrator account permissions, the system grants read-only rights for any permissions you do not assign to the cluster administrator.

Related reference

[LDAP details](#) on page 39

The LDAP page on the Cluster tab provides information about the following settings.

Enabling LDAP authentication with the Element API

You can configure storage system integration with an existing LDAP server. This enables LDAP administrators to centrally manage storage system access for users.

About this task

You can configure LDAP with either the Element user interface or the Element API. This procedure describes how to configure LDAP using the Element API.

To leverage LDAP authentication on a SolidFire cluster, you enable LDAP authentication first on the cluster using the `EnableLdapAuthentication` API method.

Steps

1. Enable LDAP authentication first on the cluster using the `EnableLdapAuthentication` API method.
2. Enter the required information.

[LDAP details](#) on page 39

```
{
  "method": "EnableLdapAuthentication",
  "params": {
    "authType": "SearchAndBind",
    "groupSearchBaseDN": "dc=prodtest,dc=solidfire,dc=net",
    "groupSearchType": "ActiveDirectory",
    "searchBindDN": "SFReadOnly@prodtest.solidfire.net",
    "searchBindPassword": "ReadOnlyPW",
    "userSearchBaseDN": "dc=prodtest,dc=solidfire,dc=net",
    "userSearchFilter": "(&(objectClass=person)(sAMAccountName=%USERNAME%))",
    "serverURIs": [
      "ldap://172.27.1.189",
    ],
  },
  "id": "1"
}
```

3. Change the values of the following parameters:

Parameters used	Description
authType: SearchAndBind	Dictates that the cluster will use the readonly service account to first search for the user being authenticated and subsequently bind that user if found and authenticated.
groupSearchBaseDN: dc=prodtest,dc=solidfire,dc=net	Specifies the location in the LDAP tree to begin searching for groups. For this example, we've used the root of our tree. If your LDAP tree is very large, you might want to set this to a more granular sub-tree to decrease search times.
userSearchBaseDN: dc=prodtest,dc=solidfire,dc=net	Specifies the location in the LDAP tree to begin searching for users. For this example, we've used the root of our tree. If your LDAP tree is very large, you might want to set this to a more granular sub-tree to decrease search times.
groupSearchType: ActiveDirectory	Uses the Windows Active Directory server as the LDAP server.

Parameters used	Description
<p>userSearchFilter:</p> <pre>"(&(objectClass=person) (sAMAccountName=%USERNAME%))"</pre> <p>If you want to use the userPrincipalName (email address for login) you could change the userSearchFilter to:</p> <pre>"(&(objectClass=person) (userPrincipalName=%USERNAME%))"</pre> <p>Or, if you'd like to search both userPrincipalName and sAMAccountName, you can use the following userSearchFilter:</p> <pre>"(&(objectClass=person)((sAMAccountName=%USERNAME%)(userPrincipalName=%USERNAME%)))"</pre>	Leverages the sAMAccountName as our username for logging in to the SolidFire cluster. These settings tell LDAP to search for the username specified during login in the sAMAccountName attribute and also limit the search to entries that have “person” as a value in the objectClass attribute.
searchBindDN	This is the distinguished name of readonly user that will be used to search the LDAP directory. For active directory it's usually easiest to use the userPrincipalName (email address format) for the user.
searchBindPassword	This is the password for the readonly user account.

After you finish

To test this, log out of the Element UI and log back in as a user in that group.

LDAP details

The LDAP page on the Cluster tab provides information about the following settings.

Note: You must enable LDAP to view these LDAP configuration settings.

Host Name/IP Address

Address of an LDAP or LDAPS directory server.

Auth Type

The user authentication method. Possible values:

- Direct Bind
- Search And Bind

Search Bind DN

A fully qualified DN to log in with to perform an LDAP search for the user (needs bind-level access to the LDAP directory).

Search Bind Password

Password used to authenticate access to the LDAP server.

User Search Base DN

The base DN of the tree used to start the user search. The system searches the subtree from the specified location.

User Search Filter

Enter the following using your domain name:

```
(&(objectClass=person)(|(sAMAccountName=%USERNAME%)(userPrincipalName=%USERNAME%)))
```

Group Search Type

Type of search that controls the default group search filter used. Possible values:

- Active Directory: Nested membership of all of a user's LDAP groups.
- No Groups: No group support.
- Member DN: Member DN-style groups (single-level).

Group Search Base DN

The base DN of the tree used to start the group search. The system searches the subtree from the specified location.

Test User Authentication

After LDAP is configured, use this to test the user name and password authentication for the LDAP server. Enter an account that already exists to test this. The distinguished name and user group information appears, which you can copy for later use when creating cluster administrators.

Testing the LDAP Configuration

After configuring LDAP, you should test it by using either the Element UI or the Element API `TestLdapAuthentication` method.

Steps

1. To test the LDAP configuration with the Element UI, do the following:
 - a. Click **Cluster > LDAP**.
 - b. Click **Test LDAP Authentication**.
 - c. Resolve any issues by using the information in the table below:

Error message	Description
<code>xLDAPUserNotFound</code>	<ul style="list-style-type: none"> • The user being tested was not found in the configured <code>userSearchBaseDN</code> subtree. • The <code>userSearchFilter</code> is configured incorrectly.
<code>xLDAPBindFailed (Error: Invalid credentials)</code>	<ul style="list-style-type: none"> • The username being tested is a valid LDAP user, but the password provided is incorrect. • The username being tested is a valid LDAP user, but the account is currently disabled.
<code>xLDAPSearchBindFailed (Error: Can't contact LDAP server)</code>	The LDAP server URI is incorrect.
<code>xLDAPSearchBindFailed (Error: Invalid credentials)</code>	The read-only username or password is configured incorrectly.
<code>xLDAPSearchFailed (Error: No such object)</code>	The <code>userSearchBaseDN</code> is not a valid location within the LDAP tree.

Error message	Description
xLDAPSearchFailed (Error: Referral)	<ul style="list-style-type: none"> The <code>userSearchBaseDN</code> is not a valid location within the LDAP tree. The <code>userSearchBaseDN</code> and <code>groupSearchBaseDN</code> are in a nested OU. This can cause permission issues. The workaround is to include the OU in the user and group base DN entries, (for example: <code>ou=storage, cn=company, cn=com</code>)

2. To test the LDAP configuration with the Element API, do the following:

- Call the `TestLdapAuthentication` method.

```
{
  "method": "TestLdapAuthentication",
  "params": {
    "username": "admin1",
    "password": "admin1PASS"
  },
  "id": 1
}
```

- Review the results. If the API call is successful, the results include the specified user's distinguished name and a list of groups in which the user is a member.

```
{
  "id": 1
  "result": {
    "groups": [
      "CN=StorageMgmt,OU=PTUsers,DC=prodtest,DC=solidfire,DC=net"
    ],
    "userDN": "CN=Admin1 Jones,OU=PTUsers,DC=prodtest,DC=solidfire,DC=net"
  }
}
```

Disabling LDAP

You can disable LDAP integration using the Element UI.

Before you begin

You have made a note of all the configuration settings, because disabling LDAP erases all settings.

Steps

- Click **Cluster > LDAP**.
- Click **No**.
- Click **Disable LDAP**.

Enabling Multi-factor authentication

Multi-factor authentication (MFA) uses a third-party Identity Provider (IdP) via the Security Assertion Markup Language (SAML) to manage user sessions. MFA enables administrators to configure additional factors of authentication as required, such as password and text message, and password and email message.

Setting up Multi-factor authentication

You can use these basic steps via the Element API to set up your cluster to use multi-factor authentication. Details of each API method can be found in the Element API Reference Guide.

Steps

1. Create a new third-party Identity Provider (IdP) configuration for the cluster by calling the following API method and passing the IdP metadata in JSON format:

CreateIdpConfiguration

IdP metadata, in plain text format, is retrieved from the third-party IdP. This metadata needs to be validated to ensure that it is correctly formatted in JSON. There are numerous JSON formatter applications available that you can use, for example:.

2. Retrieve cluster metadata, via `spMetadataUrl`, to copy to the third-party IdP by calling the following API method:

ListIdpConfigurations

`spMetadataUrl` is a URL used to retrieve service provider metadata from the cluster for the IdP in order to establish a trust relationship.

3. Configure SAML assertions on the third-party IdP to include the “NameID” attribute to uniquely identify a user for audit logging and for Single Logout to function properly.
4. Create one or more cluster administrator user accounts authenticated by a third-party IdP for authorization by calling the following API method:

AddIdpClusterAdmin

Note: The username for the IdP cluster Administrator should match the SAML attribute Name/Value mapping for the desired effect, as shown in the following examples:

- `email=bob@company.com` – where the IdP is configured to release an email address in the SAML attributes.
- `group=cluster-administrator` - where the IdP is configured to release a group property in which all users should have access.

Note too that the SAML attribute Name/Value pairing is case-sensitive for security purposes.

5. Enable MFA for the cluster by calling the following API method:

EnableIdpAuthentication

Additional information for Multi-factor authentication

You should be aware of the following caveats in relation to multi-factor authentication.

- In order to refresh IdP certificates that are no longer valid, you will need to use a non-IdP admin user to call the following API method:

UpdateIdpConfiguration

- MFA is incompatible with certificates that are less than 2048 bits in length. By default, a 2048-bit SSL certificate is created on the cluster. You should avoid setting a smaller sized certificate when calling the API method:

SetSSLCertificate

Note: If the cluster is using a certificate that is less than 2048 bits pre-upgrade, the cluster certificate must be updated with a 2048-bit or greater certificate after upgrade to Element 12.0 or later.

- IdP admin users cannot be used to make API calls directly (for example, via SDKs or Postman) or used for other integrations (for example, OpenStack Cinder or vCenter Plug-in). Add either

LDAP cluster admin users or local cluster admin users if you need to create users that have these abilities.

Related information

[Managing storage with the Element API](#)

Configuring cluster settings

You can view and change cluster-wide settings and perform cluster-specific tasks from the Cluster tab of the Element UI.

You can configure settings such as cluster fullness threshold, support access, encryption at rest, virtual volumes, SnapMirror, and NTP broadcast client.

Related concepts

[Working with virtual volumes](#) on page 80

You can view information and perform tasks for virtual volumes and their associated storage containers, protocol endpoints, bindings, and hosts using the Element UI.

[Managing the Terms of Use banner](#) on page 45

You can configure a banner that contains a message for the user.

[Managing SNMP](#) on page 46

You can configure Simple Network Management Protocol (SNMP) in your cluster.

[Managing drives](#) on page 48

Each node contains one or more physical drives that are used to store a portion of the data for the cluster. The cluster utilizes the capacity and performance of the drive after the drive has been successfully added to a cluster. You can use the Element UI to manage drives.

[Managing nodes](#) on page 49

You can manage SolidFire storage and Fibre Channel nodes from the Nodes page of the Cluster tab.

[Managing virtual networks](#) on page 54

Virtual networking in SolidFire storage enables traffic between multiple clients that are on separate logical networks to be connected to one cluster. Connections to the cluster are segregated in the networking stack through the use of VLAN tagging.

Related tasks

[Using SnapMirror replication between Element and ONTAP clusters](#) on page 124

You can create SnapMirror relationships from the Data Protection tab in the NetApp Element UI. SnapMirror functionality must be enabled to see this in the user interface.

[Setting cluster full threshold](#) on page 44

You can change the level at which the system generates a block cluster fullness warning using the steps below. In addition, you can use the `ModifyClusterFullThreshold` API method to change the level at which the system generates a block or metadata warning.

[Enabling and disabling support access](#) on page 44

You can enable support access to temporarily allow NetApp support personnel access to storage nodes via SSH for troubleshooting.

[Enabling and disabling encryption for a cluster](#) on page 44

You can enable and disable cluster-wide encryption at rest. This feature is not enabled by default.

[Enabling a broadcast client](#) on page 46

You can instruct each node in a cluster to listen for Network time protocol (NTP) broadcasts instead of querying an NTP server for updates by using the broadcast client setting.

[Viewing Fibre Channel ports details](#) on page 53

You can view details of Fibre Channel ports such as its status, name, and port address from the FC Ports page.

Related information

[How are the blockSpace thresholds calculated for Element](#)

Setting cluster full threshold

You can change the level at which the system generates a block cluster fullness warning using the steps below. In addition, you can use the `ModifyClusterFullThreshold` API method to change the level at which the system generates a block or metadata warning.

Before you begin

You have cluster administrator privileges.

Steps

1. Click **Cluster > Settings**.
2. In the Cluster Full Settings section, enter a percentage in **Raise a warning alert when _% capacity remains before Helix could not recover from a node failure**.
3. Click **Save Changes**.

Related information

[How are the blockSpace thresholds calculated for Element](#)

Enabling and disabling support access

You can enable support access to temporarily allow NetApp support personnel access to storage nodes via SSH for troubleshooting.

Before you begin

You must have cluster admin privileges to change support access.

Steps

1. Click **Cluster > Settings**.
2. In the Enable / Disable Support Access section, enter the duration (in hours) that you want to allow support to have access.
3. Click **Enable Support Access**.
4. Optional: To disable support access, click **Disable Support Access**.

Enabling and disabling encryption for a cluster

You can enable and disable cluster-wide encryption at rest. This feature is not enabled by default.

Before you begin

- You must have cluster administrator privileges to change encryption settings.
- Ensure that the cluster is in a healthy state before changing encryption settings.

Tip: Configure NTP on the cluster to point to a local NTP server. You should use the IP address and not the DNS host name. The default NTP server at cluster creation time is set to `us.pool.ntp.org`; however, a connection to this site cannot always be made depending on the physical location of the SolidFire cluster.

Steps

1. Click **Cluster > Settings**.
2. Click **Enable Encryption at Rest**.

3. Optional: To disable encryption at rest, click **Disable Encryption at Rest**.

Related concepts

[Encryption at rest](#) on page 45

SolidFire clusters enable you to encrypt all data stored on the cluster.

Encryption at rest

SolidFire clusters enable you to encrypt all data stored on the cluster.

All drives in storage nodes capable of encryption leverage AES 256-bit encryption at the drive level. Each drive has its own encryption key, which is created when the drive is first initialized. When you enable the encryption feature, a cluster-wide password is created, and chunks of the password are then distributed to all nodes in the cluster. No single node stores the entire password. The password is then used to password-protect all access to the drives. The password is needed to unlock the drive and then not needed unless power is removed from the drive or the drive is locked.

Enabling the encryption at rest feature does not affect performance or efficiency on the cluster. Additionally, if an encryption-enabled drive or node is removed from the cluster with the Element API or Element UI, encryption at rest will be disabled on the drives. After the drive is removed, the drive can be securely erased by using the `SecureEraseDrives` API method. If a drive or node is forcibly removed from the cluster, the data remains protected by the cluster-wide password and the drive's individual encryption keys.

Managing the Terms of Use banner

You can configure a banner that contains a message for the user.

Enabling Terms of Use

You can enable a Terms of Use banner that appears when a user logs in to the Element UI. When the user clicks on the banner, a text dialog box appears containing the message you have configured for the cluster. The banner can be dismissed at any time.

Before you begin

You must have cluster administrator privileges to enable Terms of Use functionality.

Steps

1. Click **Users > Terms of Use**.
2. In the **Terms of Use** form, enter the text to be displayed for the Terms of Use dialog box.

Note: Do not exceed 4096 characters.

3. Click **Enable**.

Editing Terms of Use

You can edit the text that a user sees when they select the Terms of Use login banner.

Before you begin

- You must have cluster administrator privileges to configure Terms of Use.
- Ensure that the Terms of Use feature is enabled.

Steps

1. Click **Users > Terms of Use**.
2. In the **Terms of Use** dialog box, edit the text that you want to appear.

Note: Do not exceed 4096 characters.

3. Click **Save Changes**.

Disabling Terms of Use

You can disable the Terms of Use banner. With the banner disabled, the user is no longer requested to accept the terms of use when using the Element UI.

Before you begin

- You must have cluster administrator privileges to configure Terms of Use.
- Ensure that Terms of Use is enabled.

Steps

1. Click **Users > Terms of Use**.
2. Click **Disable**.

Enabling a broadcast client

You can instruct each node in a cluster to listen for Network time protocol (NTP) broadcasts instead of querying an NTP server for updates by using the broadcast client setting.

Before you begin

- You must have cluster administrator privileges to configure this setting.
- You must configure an NTP server on your network as a broadcast server.

About this task

The NTP is used to synchronize clocks over a network. Connection to an internal or external NTP server should be part of the initial cluster setup.

You can enter up to five different NTP servers.

Note: You can use both IPv4 and IPv6 addresses.

Steps

1. Click **Cluster > Settings**.
2. Under Network Time Protocol Settings, select **Yes** to use as a broadcast client.
3. In the **Server** field, enter the NTP server you configured in broadcast mode.
4. Click **Save Changes**.

Managing SNMP

You can configure Simple Network Management Protocol (SNMP) in your cluster.

You can select an SNMP requestor, select which version of SNMP to use, identify the SNMP User-based Security Model (USM) user, and configure traps to monitor the SolidFire cluster. You can also view and access management information base files.

Note: You can use both IPv4 and IPv6 addresses.

SNMP details

On the SNMP page of the Cluster tab, you can view the following information.

SNMP MIBs

The MIB files that are available for you to view or download.

General SNMP Settings

You can enable or disable SNMP. After you enable SNMP, you can choose which version to use. If using version 2, you can add requestors, and if using version 3, you can set up USM users.

SNMP Trap Settings

You can identify which traps you want to capture. You can set the host, port, and community string for each trap recipient.

Configuring an SNMP requestor

When SNMP version 2 is enabled, you can enable or disable a requestor, and configure requestors to receive authorized SNMP requests.

Steps

1. Click **Cluster > SNMP**.
2. Under **General SNMP Settings**, click **Yes** to enable SNMP.
3. From the **Version** list, select **Version 2**.
4. In the **Requestors** section, enter the **Community String** and **Network** information.

Note: By default, the community string is public, and the network is localhost. You can change these default settings.

5. Optional: To add another requestor, click **Add a Requestor** and enter the **Community String** and **Network** information.
6. Click **Save Changes**.

Related tasks

[Configuring SNMP traps](#) on page 48

System administrators can use SNMP traps, also referred to as notifications, to monitor the health of the SolidFire cluster.

[Viewing managed object data using management information base files](#) on page 48

You can view and download the management information base (MIB) files used to define each of the managed objects. The SNMP feature supports read-only access to those objects defined in the SolidFire-StorageCluster-MIB.

Configuring an SNMP USM user

When you enable SNMP version 3, you need to configure a USM user to receive authorized SNMP requests.

Steps

1. Click **Cluster > SNMP**.
2. Under **General SNMP Settings**, click **Yes** to enable SNMP.
3. From the **Version** list, select **Version 3**.
4. In the **USM Users** section, enter the name, password, and passphrase.
5. Optional: To add another USM user, click **Add a USM User** and enter the name, password, and passphrase.
6. Click **Save Changes**.

Configuring SNMP traps

System administrators can use SNMP traps, also referred to as notifications, to monitor the health of the SolidFire cluster.

About this task

When SNMP traps are enabled, the SolidFire cluster generates traps associated with event log entries and system alerts. To receive SNMP notifications, you need to choose the traps that should be generated and identify the recipients of the trap information. By default, no traps are generated.

Steps

1. Click **Cluster > SNMP**.
2. Select one or more types of traps in the **SNMP Trap Settings** section that the system should generate:
 - Cluster Fault Traps
 - Cluster Resolved Fault Traps
 - Cluster Event Traps
3. In the **Trap Recipients** section, enter the host, port, and community string information for a recipient.
4. Optional: To add another trap recipient, click **Add a Trap Recipient** and enter host, port, and community string information.
5. Click **Save Changes**.

Viewing managed object data using management information base files

You can view and download the management information base (MIB) files used to define each of the managed objects. The SNMP feature supports read-only access to those objects defined in the SolidFire-StorageCluster-MIB.

About this task

The statistical data provided in the MIB shows system activity for the following:

- Cluster statistics
- Volume statistics
- Volumes by account statistics
- Node statistics
- Other data such as reports, errors, and system events

The system also supports access to the MIB file containing the upper level access points (OIDS) to SF-Series products.

Steps

1. Click **Cluster > SNMP**.
2. Under **SNMP MIBs**, click the MIB file you want to download.
3. In the resulting download window, open or save the MIB file.

Managing drives

Each node contains one or more physical drives that are used to store a portion of the data for the cluster. The cluster utilizes the capacity and performance of the drive after the drive has been successfully added to a cluster. You can use the Element UI to manage drives.

Related tasks

[Adding drives to a cluster](#) on page 20

When you add a node to the cluster or install new drives in an existing node, the drives automatically register as available. You must add the drives to the cluster by using either the Element UI or API before they can participate in the cluster.

Drives details

The Drives page on the Cluster tab provides a list of the active drives in the cluster. You can filter the page by selecting from the **Active**, **Available**, **Removing**, **Erasing**, and **Failed** tabs.

When you first initialize a cluster, the active drives list is empty. You can add drives that are unassigned to a cluster and listed in the **Available** tab after a new SolidFire cluster is created.

The following elements appear in the list of active drives.

Drive ID

The sequential number assigned to the drive.

Node ID

The node number assigned when the node is added to the cluster.

Node Name

The name of the node that houses the drive.

Slot

The slot number where the drive is physically located.

Capacity

The size of the drive, in GB.

Serial

The serial number of the drive.

Wear Remaining

The wear level indicator.

The storage system reports the approximate amount of wear available on each solid-state drive (SSD) for writing and erasing data. A drive that has consumed 5 percent of its designed write and erase cycles reports 95 percent wear remaining. The system does not refresh drive wear information automatically; you can refresh or close and reload the page to refresh the information.

Type

The type of drive. The type can be either block or metadata.

Managing nodes

You can manage SolidFire storage and Fibre Channel nodes from the Nodes page of the Cluster tab.

If a newly added node accounts for more than 50 percent of the total cluster capacity, some of the capacity of this node is made unusable ("stranded"), so that it complies with the capacity rule. This remains the case until more storage is added. If a very large node is added that also disobeys the capacity rule, the previously stranded node will no longer be stranded, while the newly added node becomes stranded. Capacity should always be added in pairs to avoid this happening. When a node becomes stranded, an appropriate cluster fault is thrown.

Related tasks

[Adding a node to a cluster](#) on page 50

You can add nodes to a cluster when more storage is needed or after cluster creation. Nodes require initial configuration when they are first powered on. After the node is configured, it appears in the list of pending nodes and you can add it to a cluster.

Related reference

[Node states](#) on page 18

A node can be in one of several states depending on the level of configuration.

Adding a node to a cluster

You can add nodes to a cluster when more storage is needed or after cluster creation. Nodes require initial configuration when they are first powered on. After the node is configured, it appears in the list of pending nodes and you can add it to a cluster.

About this task

The software version on each node in a cluster must be compatible. When you add a node to a cluster, the cluster installs the cluster version of Element software on the new node as needed.

You can add nodes of smaller or larger capacities to an existing cluster. You can add larger node capacities to a cluster to allow for capacity growth. Larger nodes added to a cluster with smaller nodes must be added in pairs. This allows for sufficient space for Double Helix to move the data should one of the larger nodes fail. You can add smaller node capacities to a larger node cluster to improve performance.

Note: If a newly added node accounts for more than 50 percent of the total cluster capacity, some of the capacity of this node is made unusable ("stranded"), so that it complies with the capacity rule. This remains the case until more storage is added. If a very large node is added that also disobeys the capacity rule, the previously stranded node will no longer be stranded, while the newly added node becomes stranded. Capacity should always be added in pairs to avoid this happening. When a node becomes stranded, the strandedCapacity cluster fault is thrown.

[NetApp video: Scale on Your Terms: Expanding a SolidFire Cluster](#)

Steps

1. Select **Cluster > Nodes**.
2. Click **Pending** to view the list of pending nodes.
3. Do one of the following:
 - To add individual nodes, click the **Actions** icon for the node you want to add.
 - To add multiple nodes, select the check box of the nodes to add, and then **Bulk Actions**.

Note: If the node you are adding has a different version of Element software than the version running on the cluster, the cluster asynchronously updates the node to the version of Element software running on the cluster master. After the node is updated, it automatically adds itself to the cluster. During this asynchronous process, the node will be in a pendingActive state.

4. Click **Add**.

The node appears in the list of active nodes.

Related concepts

[Node versioning and compatibility](#) on page 51

Node compatibility is based on the Element software version installed on a node. Element software-based storage clusters automatically image a node to the Element software version on the cluster if the node and cluster are not at compatible versions.

[Configuring a Fibre Channel node](#) on page 21

Fibre Channel nodes enable you to connect the cluster to a Fibre Channel network fabric. Fibre Channel nodes are added in pairs, and operate in active-active mode (all nodes actively process traffic for the cluster). Clusters running Element software version 9.0 and later support up to four nodes; clusters running previous versions support a maximum of two nodes.

Node versioning and compatibility

Node compatibility is based on the Element software version installed on a node. Element software-based storage clusters automatically image a node to the Element software version on the cluster if the node and cluster are not at compatible versions.

The following list describes the software release significance levels that make up the Element software version number:

Major

The first number designates a software release. A node with one major component number cannot be added to a cluster containing nodes of a different major-patch number, nor can a cluster be created with nodes of mixed major versions.

Minor

The second number designates smaller software features or enhancements to existing software features that have been added to a major release. This component is incremented within a major version component to indicate that this incremental release is not compatible with any other Element software incremental releases with a different minor component.

For example, 11.0 is not compatible with 11.1, and 11.1 is not compatible with 11.2.

Micro

The third number designates a compatible patch (incremental release) to the Element software version represented by the major.minor components. For example, 11.0.1 is compatible with 11.0.2, and 11.0.2 is compatible with 11.0.3.

Major and minor version numbers must match for compatibility. Micro numbers do not have to match for compatibility.

Cluster capacity in a mixed node environment

You can mix different types of nodes in a cluster. The SF-Series 2405, 3010, 4805, 6010, 9605, 9010, 19210, 38410 and the H-Series can coexist in a cluster.

The H-Series consists of H610S-1, H610S-2, H610S-4, and H410S nodes. These nodes are both 10GbE and 25GbE capable.

It is best to not intermix non-encrypted and encrypted nodes. In a mixed node cluster, no node can be larger than 33 percent of the total cluster capacity. For instance, in a cluster with four SF-Series 4805 nodes, the largest node that can be added alone is an SF-Series 9605. The cluster capacity threshold is calculated based on the potential loss of the largest node in this situation.

Beginning with Element 12.0, the following SF-series storage nodes are not supported:

- SF3010
- SF6010
- SF9010

If you upgrade one of these storage nodes to Element 12.0, you will see an error stating that this node is not supported by Element 12.0.

Node states

A node can be in one of several states depending on the level of configuration.

Available

The node has no associated cluster name and is not yet part of a cluster.

Pending

The node is configured and can be added to a designated cluster.

Authentication is not required to access the node.

Pending Active

The system is in the process of installing compatible Element software on the node. When complete, the node will move to the **Active** state.

Active

The node is participating in a cluster.

Authentication is required to modify the node.

In each of these states, some fields are read only.

Node details

On the Nodes page of the Cluster tab, you can view information about nodes such as ID, name, configured IOPs, and role type.

Node ID

The system-generated ID for the node.

Node Name

The host name for the node.

Available 4k IOPS

The IOPS configured for the node.

Node Role

The role that the node has in the cluster. Possible values:

- Cluster Master: The node that performs cluster-wide administrative tasks and contains the MVIP and SVIP.
- Ensemble Node: A node that participates in the cluster. There are either 3 or 5 ensemble nodes depending on cluster size.
- Fibre Channel: A node in the cluster.

Node Type

The model type of the node.

Active Drives

The number of active drives in the node.

Management IP

The management IP (MIP) address assigned to node for 1GbE or 10GbE network admin tasks.

Cluster IP

The cluster IP (CIP) address assigned to the node used for the communication between nodes in the same cluster.

Storage IP

The storage IP (SIP) address assigned to the node used for iSCSI network discovery and all data network traffic.

Management VLAN ID

The virtual ID for the management local area network.

Storage VLAN ID

The virtual ID for the storage local area network.

Version

The version of software running on each node.

Replication Port

The port used on nodes for remote replication.

Service Tag

The unique service tag number assigned to the node.

Viewing individual node details

You can view details for individual nodes such as service tags, drive details, and graphics for utilization and drive statistics. The Nodes page of the Cluster tab provides the Version column where you can view the software version of each node.

Steps

1. Click **Cluster > Nodes**.
2. Click the Actions icon for a node.
3. Click **View Details**.

Viewing Fibre Channel ports details

You can view details of Fibre Channel ports such as its status, name, and port address from the FC Ports page.

Steps

1. Click **Cluster > FC Ports**.
2. To filter information on this page, click **Filter**.

Related reference

[Fibre Channel ports details](#) on page 53

The FC Ports page on the Cluster tab provides information about the Fibre Channel ports that are connected to the cluster.

Fibre Channel ports details

The FC Ports page on the Cluster tab provides information about the Fibre Channel ports that are connected to the cluster.

The following list describes information about the Fibre Channel ports that are connected to the cluster:

Node ID

The node hosting the session for the connection.

Node Name

System-generated node name.

Slot

Slot number where the Fibre Channel port is located.

HBA Port

Physical port on the Fibre Channel host bus adapter (HBA).

WWNN

The world wide node name.

WWPN

The target world wide port name.

Switch WWN

World wide name of the Fibre Channel switch.

Port State

Current state of the port.

nPort ID

The node port ID on the Fibre Channel fabric.

Speed

The negotiated Fibre Channel speed. Possible values are as follows:

- 4Gbps
- 8Gbps
- 16Gbps

Managing virtual networks

Virtual networking in SolidFire storage enables traffic between multiple clients that are on separate logical networks to be connected to one cluster. Connections to the cluster are segregated in the networking stack through the use of VLAN tagging.

Related tasks

[Adding a virtual network](#) on page 55

You can add a new virtual network to a cluster configuration to enable a multi-tenant environment connection to a cluster running Element software.

[Enabling virtual routing and forwarding](#) on page 56

You can enable virtual routing and forwarding (VRF), which allows multiple instances of a routing table to exist in a router and work simultaneously. This functionality is available for storage networks only.

[Editing a virtual network](#) on page 56

You can change VLAN attributes, such as VLAN name, netmask, and size of the IP address blocks. The VLAN tag and SVIP cannot be modified for a VLAN. The gateway attribute is not a valid parameter for non-VRF VLANs.

[Editing VRF VLANs](#) on page 56

You can change VRF VLAN attributes, such as VLAN name, netmask, gateway, and IP address blocks.

[Deleting a virtual network](#) on page 57

You can remove a virtual network object. You must add the address blocks to another virtual network before you remove a virtual network.

Virtual networks details

On the Network page of the Cluster tab, you can view information about virtual networks, such as ID, VLAN Tag, SVIP, and Netmask.

ID

Unique ID of the VLAN network, which is assigned by the system.

Name

Unique user-assigned name for the VLAN network.

VLAN Tag

VLAN tag assigned when the virtual network was created.

SVIP

Storage virtual IP address assigned to the virtual network.

Netmask

Netmask for this virtual network.

Gateway

Unique IP address of a virtual network gateway. VRF must be enabled.

VRF Enabled

Indication of whether virtual routing and forwarding is enabled or not.

IPs Used

The range of virtual network IP addresses used for the virtual network.

Adding a virtual network

You can add a new virtual network to a cluster configuration to enable a multi-tenant environment connection to a cluster running Element software.

Before you begin

- Identify the block of IP addresses that will be assigned to the virtual networks on the cluster nodes.
- Identify a storage network IP (SVIP) address that will be used as an endpoint for all NetApp Element storage traffic.



Attention: You must consider the following criteria for this configuration:

- VLANs that are not VRF-enabled require initiators to be in the same subnet as the SVIP.
- VLANs that are VRF-enabled do not require initiators to be in the same subnet as the SVIP, and routing is supported.
- The default SVIP does not require initiators to be in the same subnet as the SVIP, and routing is supported.

About this task

When a virtual network is added, an interface for each node is created and each requires a virtual network IP address. The number of IP addresses you specify when creating a new virtual network must be equal to or greater than the number of nodes in the cluster. Virtual network addresses are bulk provisioned by and assigned to individual nodes automatically. You do not need to manually assign virtual network addresses to the nodes in the cluster.

Steps

- Click **Cluster > Network**.
- Click **Create VLAN**.
- In the **Create a New VLAN** dialog box, enter values in the following fields:
 - VLAN Name**
 - VLAN Tag**
 - SVIP**
 - Netmask**
 - (Optional) **Description**

4. Enter the **Starting IP** address for the range of IP addresses in **IP Address Blocks**.
5. Enter the **Size** of the IP range as the number of IP addresses to include in the block.
6. Click **Add a Block** to add a non-continuous block of IP addresses for this VLAN.
7. Click **Create VLAN**.

Enabling virtual routing and forwarding

You can enable virtual routing and forwarding (VRF), which allows multiple instances of a routing table to exist in a router and work simultaneously. This functionality is available for storage networks only.

About this task

You can enable VRF only at the time of creating a VLAN. If you want to switch back to non-VRF, you must delete and re-create the VLAN.

Steps

1. Click **Cluster > Network**.
2. To enable VRF on a new VLAN, select **Create VLAN**.
 - a. Enter relevant information for the new VRF/VLAN. See [Adding a virtual network](#).
 - b. Select the **Enable VRF** check box.
 - c. Optional: Enter a gateway.
3. Click **Create VLAN**.

Related tasks

[Adding a virtual network](#) on page 55

You can add a new virtual network to a cluster configuration to enable a multi-tenant environment connection to a cluster running Element software.

Editing a virtual network

You can change VLAN attributes, such as VLAN name, netmask, and size of the IP address blocks. The VLAN tag and SVIP cannot be modified for a VLAN. The gateway attribute is not a valid parameter for non-VRF VLANs.

About this task

If any iSCSI, remote replication, or other network sessions exist, the modification might fail.

Steps

1. Click **Cluster > Network**.
2. Click the Actions icon for the VLAN you want to edit.
3. Click **Edit**.
4. In the **Edit VLAN** dialog box, enter the new attributes for the VLAN.
5. Click **Add a Block** to add a non-continuous block of IP addresses for the virtual network.
6. Click **Save Changes**.

Editing VRF VLANs

You can change VRF VLAN attributes, such as VLAN name, netmask, gateway, and IP address blocks.

Steps

1. Click **Cluster > Network**.
2. Click the Actions icon for the VLAN you want to edit.
3. Click **Edit**.

4. Enter the new attributes for the VRF VLAN in the **Edit VLAN** dialog box.
5. Click **Save Changes**.

Deleting a virtual network

You can remove a virtual network object. You must add the address blocks to another virtual network before you remove a virtual network.

Steps

1. Click **Cluster > Network**.
2. Click the Actions icon for the VLAN you want to delete.
3. Click **Delete**.
4. Confirm the message.

Related tasks

[Editing a virtual network](#) on page 56

You can change VLAN attributes, such as VLAN name, netmask, and size of the IP address blocks. The VLAN tag and SVIP cannot be modified for a VLAN. The gateway attribute is not a valid parameter for non-VRF VLANs.

Creating a cluster supporting FIPS drives

Security is becoming increasingly critical for the deployment of solutions in many customer environments. Federal Information Processing Standards (FIPS) are standards for computer security and interoperability. FIPS 140-2 certified encryption for data at rest is a component of the overall security solution.

Steps

1. [Avoiding mixing nodes for FIPS drives](#) on page 58
To prepare for enabling the FIPS drives feature, you should avoid mixing nodes where some are FIPS drives capable and some are not.
2. [Enabling encryption at rest](#) on page 58
You can enable and disable cluster-wide encryption at rest. This feature is not enabled by default. To support FIPS drives, you must enable encryption at rest.
3. [Identifying whether nodes are ready for the FIPS drives feature](#) on page 58
You should check to see if all nodes in the storage cluster are ready to support FIPS drives by using the NetApp Element software GetFipsReport API method.
4. [Enabling the FIPS drives feature](#) on page 59
You can enable the FIPS drives feature by using the NetApp Element software EnableFeature API method.
5. [Checking the FIPS drive status](#) on page 59
You can check whether the FIPS drives feature is enabled on the cluster by using the NetApp Element software GetFeatureStatus API method, which shows whether the FIPS Drives Enabled Status is true or false.
6. [Troubleshooting the FIPS drive feature](#) on page 59

Using the NetApp Element software UI, you can view alerts for information about cluster faults or errors in the system related to the FIPS drives feature.

Avoiding mixing nodes for FIPS drives

To prepare for enabling the FIPS drives feature, you should avoid mixing nodes where some are FIPS drives capable and some are not.

About this task

A cluster is considered FIPS drives compliant based on the following conditions:

- All drives are certified as FIPS drives.
- All nodes are FIPS drives nodes.
- Encryption at Rest (EAR) is enabled.
- The FIPS drives feature is enabled. All drives and nodes must be FIPS capable and Encryption at Rest must be enabled in order to enable the FIPS drive feature.

Enabling encryption at rest

You can enable and disable cluster-wide encryption at rest. This feature is not enabled by default.

To support FIPS drives, you must enable encryption at rest.

Steps

1. In the NetApp Element software UI, click **Cluster > Settings**.
2. Click **Enable Encryption at Rest**.

Related concepts

[Encryption at rest](#) on page 45

SolidFire clusters enable you to encrypt all data stored on the cluster.

Related tasks

[Enabling and disabling encryption for a cluster](#) on page 44

You can enable and disable cluster-wide encryption at rest. This feature is not enabled by default.

Identifying whether nodes are ready for the FIPS drives feature

You should check to see if all nodes in the storage cluster are ready to support FIPS drives by using the NetApp Element software `GetFipsReport` API method.

About this task

The resulting report shows one of the following statuses:

- None: Node is not capable of supporting the FIPS drives feature.
- Partial: Node is FIPS capable, but not all drives are FIPS drives.
- Ready: Node is FIPS capable and all drives are FIPS drives or no drives are present.

Steps

1. Using the Element API, check to see if the nodes and drives in the storage cluster are capable of FIPS drives by entering:
`GetFipsReport`
2. Review the results, noting any nodes that did not display a status of Ready.
3. For any nodes that did not display a Ready status, check to see if the drive is capable of supporting the FIPS drives feature:
 - Using the Element API, enter: `GetHardwareList`

- Note the value of the **DriveEncryptionCapabilityType**. If it is "fips," the hardware can support the FIPS drives feature.

See details about `GetFipsReport` or `ListDriveHardware` in the *NetApp Element API Reference Guide*.

4. If the drive cannot support the FIPS drives feature, replace the hardware with FIPS hardware (either node or drives).

Related information

[Managing storage with the Element API](#)

Enabling the FIPS drives feature

You can enable the FIPS drives feature by using the NetApp Element software `EnableFeature` API method.

Before you begin

Encryption at Rest must be enabled on the cluster and all nodes and drives must be FIPS capable, as indicated when the `GetFipsReport` displays a Ready status for all nodes.

Step

Using the Element API, enable FIPS on all drives by entering:

`EnableFeature params: FipsDrives`

Related information

[Managing storage with the Element API](#)

Checking the FIPS drive status

You can check whether the FIPS drives feature is enabled on the cluster by using the NetApp Element software `GetFeatureStatus` API method, which shows whether the FIPS Drives Enabled Status is true or false.

Steps

1. Using the Element API, check the FIPS drives feature on the cluster by entering:

`GetFeatureStatus`

2. Review the results of the `GetFeatureStatus` API call. If the FIPS Drives enabled value is True, the FIPS drives feature is enabled.

```
{"enabled": true,  
"feature": "FipsDrives"  
}
```

Related information

[Managing storage with the Element API](#)

Troubleshooting the FIPS drive feature

Using the NetApp Element software UI, you can view alerts for information about cluster faults or errors in the system related to the FIPS drives feature.

Steps

1. Using the Element UI, select **Reporting > Alerts**.
2. Look for cluster faults including:
 - FIPS drives mismatched

- FIPS drives out of compliance

3. For resolution suggestions, see Cluster Fault code information.

Related reference

[Cluster fault codes](#) on page 145

The system reports an error or a state that might be of interest by generating a fault code, which is listed on the Alerts page. These codes help you determine what component of the system experienced the alert and why the alert was generated.

Enabling FIPS 140-2 for HTTPS on your cluster

You can use the `EnableFeature` API method to enable the FIPS 140-2 operating mode for HTTPS communications.

About this task

With NetApp Element software, you can choose to enable Federal Information Processing Standards (FIPS) 140-2 operating mode on your cluster. Enabling this mode activates the NetApp Cryptographic Security Module (NCSM) and leverages FIPS 140-2 Level 1 certified encryption for all communication via HTTPS to the NetApp Element UI and API.



Attention: After you enable FIPS 140-2 mode, it cannot be disabled. When FIPS 140-2 mode is enabled, each node in the cluster reboots and runs through a self-test ensuring that the NCSM is correctly enabled and operating in the FIPS 140-2 certified mode. This causes an interruption to both management and storage connections on the cluster. You should plan carefully and only enable this mode if your environment needs the encryption mechanism it offers.

For more information, see the Element API information.

The following is an example of the API request to enable FIPS:

```
{  
  "method": "EnableFeature",  
  "params": {  
    "feature" : "fips"  
  },  
  "id": 1  
}
```

After this operating mode is enabled, all HTTPS communication uses the FIPS 140-2 approved ciphers.

Related reference

[SSL ciphers](#) on page 61

SSL ciphers are encryption algorithms used by hosts to establish a secure communication. There are standard ciphers that Element software supports and non-standard ones when FIPS 140-2 mode is enabled.

Related information

[Managing storage with the Element API](#)

SSL ciphers

SSL ciphers are encryption algorithms used by hosts to establish a secure communication. There are standard ciphers that Element software supports and non-standard ones when FIPS 140-2 mode is enabled.

The following lists provide the standard Secure Socket Layer (SSL) ciphers supported by Element software and the SSL ciphers supported when FIPS 140-2 mode is enabled:

FIPS 140-2 disabled

TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048) - C
TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048) - A
TLS_RSA_WITH_AES_128_CBC_SHA256 (rsa 2048) - A
TLS_RSA_WITH_AES_128_GCM_SHA256 (rsa 2048) - A
TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048) - A
TLS_RSA_WITH_AES_256_CBC_SHA256 (rsa 2048) - A
TLS_RSA_WITH_AES_256_GCM_SHA384 (rsa 2048) - A
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA (rsa 2048) - A
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA (rsa 2048) - A
TLS_RSA_WITH_IDEA_CBC_SHA (rsa 2048) - A
TLS_RSA_WITH_RC4_128_MD5 (rsa 2048) - C
TLS_RSA_WITH_RC4_128_SHA (rsa 2048) - C
TLS_RSA_WITH_SEED_CBC_SHA (rsa 2048) - A

FIPS 140-2 enabled

TLS_DHE_RSA_WITH_AES_128_CBC_SHA256 (dh 2048) - A
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256 (dh 2048) - A
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256 (dh 2048) - A
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384 (dh 2048) - A
TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (secp256r1) - A
TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (secp256r1) - A
TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (secp256r1) - A
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (secp256r1) - A
TLS_RSA_WITH_3DES_EDE_CBC_SHA (rsa 2048) - C
TLS_RSA_WITH_AES_128_CBC_SHA (rsa 2048) - A
TLS_RSA_WITH_AES_128_CBC_SHA256 (rsa 2048) - A
TLS_RSA_WITH_AES_128_GCM_SHA256 (rsa 2048) - A
TLS_RSA_WITH_AES_256_CBC_SHA (rsa 2048) - A
TLS_RSA_WITH_AES_256_CBC_SHA256 (rsa 2048) - A
TLS_RSA_WITH_AES_256_GCM_SHA384 (rsa 2048) - A

Related tasks

[Enabling FIPS 140-2 for HTTPS on your cluster](#) on page 60

You can use the `EnableFeature` API method to enable the FIPS 140-2 operating mode for HTTPS communications.

Getting started with External key management

External key management (EKM) provides secure Authentication Key (AK) management in conjunction with an off-cluster external key server (EKS). The EKS provides secure generation and storage of the AKs.

The AKs are used to lock and unlock Self Encrypting Drives (SEDs) when Encryption At Rest (EAR) is enabled on the cluster. The cluster utilizes the Key Management Interoperability Protocol (KMIP), an OASIS defined standard protocol, to communicate with the EKS.

Steps

1. *Setting up External key management* on page 62

You can use these basic steps via the Element API to setup your external key management feature. Details of each API method can be found in the Element API Reference Guide.

2. *Recovering inaccessible or invalid Authentication Keys* on page 63

Occasionally, an error can occur that requires user intervention. In the event of an error, a cluster fault (referred to as a *cluster fault code*) will be generated. The two most likely cases are described here.

3. *External Key Management API Commands* on page 63

List of all of the APIs available for managing and configuring EKM.

Setting up External key management

You can use these basic steps via the Element API to setup your external key management feature. Details of each API method can be found in the Element API Reference Guide.

Steps

1. Establish a trust relationship with the External Key Server (EKS).

- Create a public/private key pair for the Element cluster that is used to establish a trust relationship with the key server by calling the following API method:

`CreatePublicPrivateKeyPair`

- Get the certificate sign request (CSR) which the Certification Authority needs to sign. The CSR enables the key server to verify that the Element cluster that will be accessing the keys is authenticated as the Element cluster. Call the following API method:

`GetClientCertificateSignRequest`

- Use the EKS/Certificate Authority to sign the retrieved CSR. See third-party documentation for more information.

2. Create a server and provider on the cluster to communicate with the EKS. A key provider defines where a key should be obtained, and a server defines the specific attributes of the EKS that will be communicated with.

- Create a key provider where the key server details will reside by calling the following API method:

`CreateKeyProviderKmip`

- Create a key server providing the signed certificate and the public key of the Certification Authority by calling the following API methods:

`CreateKeyServerKmip`

`TestKeyServerKmip`

If the test fails, verify your server connectivity and configuration. Then repeat the test.

- c. Add the key server into the key provider container by calling the following API methods:

`AddKeyServerToProviderKmip`

`TestKeyProviderKmip`

If the test fails, verify your server connectivity and configuration. Then repeat the test.

3. Enable encryption at rest.

- a. Enable encryption at rest by providing the ID of the key provider that contains the key server used for storing the keys by calling the following API method:

`EnableEncryptionAtRest`

Note: To enable encryption at rest using an external key management configuration, you must enable encryption at rest via the API. Enabling using the existing Element UI button will revert to using internally generated keys.

Related concepts

[Encryption at rest](#) on page 45

SolidFire clusters enable you to encrypt all data stored on the cluster.

Related tasks

[Enabling and disabling encryption for a cluster](#) on page 44

You can enable and disable cluster-wide encryption at rest. This feature is not enabled by default.

Related information

[Managing storage with the Element API](#)

Recovering inaccessible or invalid Authentication Keys

Occasionally, an error can occur that requires user intervention. In the event of an error, a cluster fault (referred to as a *cluster fault code*) will be generated. The two most likely cases are described here.

1. The cluster is unable to unlock the drives due to a `KmipServerFault` cluster fault. This can occur when the cluster first boots up and the key server is inaccessible or the required key is unavailable.
 - a. Follow the recovery steps in the cluster fault codes (if any).
2. A `sliceServiceUnhealthy` fault might be set because the metadata drives have been marked as failed and placed into the "Available" state.
 - a. Re-add the drives.
 - b. After 3 to 4 minutes, check that the `sliceServiceUnhealthy` fault has cleared.

See cluster fault code information.

[Cluster fault codes](#) on page 145

External Key Management API Commands

List of all of the APIs available for managing and configuring EKM.

Used for establishing a trust relationship between the cluster and external customer-owned servers:

- `CreatePublicPrivateKeyPair`
- `GetClientCertificateSignRequest`

Used for defining the specific details of external customer-owned servers:

- `CreateKeyServerKmip`
- `ModifyKeyServerKmip`

- `DeleteKeyServerKmip`
- `GetKeyServerKmip`
- `ListKeyServersKmip`
- `TestKeyServerKmip`

Used for creating and maintaining key providers which manage external key servers:

- `CreateKeyProviderKmip`
- `DeleteKeyProviderKmip`
- `AddKeyServerToProviderKmip`
- `RemoveKeyServerFromProviderKmip`
- `GetKeyProviderKmip`
- `ListKeyProvidersKmip`
- `TestKeyProviderKmip`

For information about the API methods, see API reference information.

[*Managing storage with the Element API*](#)

Data management

You can manage the data in a cluster running Element software from the Management tab in the Element UI. Available cluster management functions include creating and managing data volumes, user accounts, volume access groups, initiators, and Quality of Service (QoS) policies.

Related concepts

[Working with virtual volumes](#) on page 80

You can view information and perform tasks for virtual volumes and their associated storage containers, protocol endpoints, bindings, and hosts using the Element UI.

[Working with volume access groups and initiators](#) on page 89

You can use iSCSI initiators or Fibre Channel initiators to access the volumes defined within volume access groups.

Related tasks

[Working with user accounts](#) on page 65

In SolidFire storage systems, user accounts enable clients to connect to volumes on a node. When you create a volume, it is assigned to a specific user account.

[Working with volumes](#) on page 68

The SolidFire system provisions storage using volumes. Volumes are block devices accessed over the network by iSCSI or Fibre Channel clients. From the Volumes page on the Management tab, you can create, modify, clone, and delete volumes on a node. You can also view statistics about volume bandwidth and I/O usage.

Working with user accounts

In SolidFire storage systems, user accounts enable clients to connect to volumes on a node. When you create a volume, it is assigned to a specific user account.

About this task

An account contains the CHAP authentication required to access the volumes assigned to it.

An account can have up to two-thousand volumes assigned to it, but a volume can belong to only one account.

Related tasks

[Creating an account](#) on page 65

You can create an account to allow access to volumes.

[Viewing individual account performance details](#) on page 66

You can view performance activity for individual accounts in a graphical format.

[Editing an account](#) on page 67

You can edit an account to change the status, change the CHAP secrets, or modify the account name.

[Deleting an account](#) on page 67

You can delete an account when it is no longer needed.

Creating an account

You can create an account to allow access to volumes.

About this task

Each account name in the system must be unique.

Steps

1. Select **Management > Accounts**.
2. Click **Create Account**.
3. Enter a **Username**.
4. In the **CHAP Settings** section, enter the following information:
 - **Initiator Secret** for CHAP node session authentication.
 - **Target Secret** for CHAP node session authentication.

Note: Leave the credential fields blank to auto-generate either password.

5. Click **Create Account**.

Account details

The Accounts page on the Management tab provides information about each account in the system, such as ID, user name, and efficiency details for the volumes assigned to each account.

ID

The system-generated ID for the account.

Username

The name given to the account when it was created.

Status

The status of the account. Possible values:

- **active**: An active account.
- **locked**: A locked account.
- **removed**: An account that has been deleted and purged.

Active Volumes

The number of active volumes assigned to the account.

Compression

The compression efficiency score for the volumes assigned to the account.

Deduplication

The deduplication efficiency score for the volumes assigned to the account.

Thin Provisioning

The thin provisioning efficiency score for the volumes assigned to the account.

Overall Efficiency

The overall efficiency score for the volumes assigned to the account.

Viewing individual account performance details

You can view performance activity for individual accounts in a graphical format.

About this task

The graph information provides I/O and throughput information for the account. The Average and Peak activity levels are shown in increments of 10-second reporting periods. These statistics include activity for all volumes assigned to the account.

Steps

1. Select **Management > Accounts**.
2. Click the Actions icon for an account.

3. Click **View Details**.

Editing an account

You can edit an account to change the status, change the CHAP secrets, or modify the account name.

About this task

Modifying CHAP settings in an account or removing initiators or volumes from an access group can cause initiators to lose access to volumes unexpectedly. To verify that volume access will not be lost unexpectedly, always log out iSCSI sessions that will be affected by an account or access group change, and verify that initiators can reconnect to volumes after any changes to initiator settings and cluster settings have been completed.



Attention: Persistent volumes that are associated with management services are assigned to a new account that is created during installation or upgrade. If you are using persistent volumes, do not modify or delete their associated account.

Steps

1. Select **Management > Accounts**.
2. Click the Actions icon for an account.
3. In the resulting menu, select **Edit**.
4. Optional: Edit the **Username**.
5. Optional: Click the **Status** drop-down list and select a different status.

Attention: Changing the status to **locked** terminates all iSCSI connections to the account, and the account is no longer accessible. Volumes associated with the account are maintained; however, the volumes are not iSCSI discoverable.

6. Optional: Under **CHAP Settings**, edit the **Initiator Secret** and **Target Secret** credentials used for node session authentication.

Note: If you do not change the **CHAP Settings** credentials, they remain the same. If you make the credentials fields blank, the system generates new passwords.

7. Click **Save Changes**.

Deleting an account

You can delete an account when it is no longer needed.

Before you begin

Delete and purge any volumes associated with the account before you delete the account.



Attention: Persistent volumes that are associated with management services are assigned to a new account that is created during installation or upgrade. If you are using persistent volumes, do not modify or delete their associated account.

Steps

1. Select **Management > Accounts**.

2. Click the Actions icon for the account you want to delete.
3. In the resulting menu, select **Delete**.
4. Confirm the action.

Working with volumes

The SolidFire system provisions storage using volumes. Volumes are block devices accessed over the network by iSCSI or Fibre Channel clients. From the Volumes page on the Management tab, you can create, modify, clone, and delete volumes on a node. You can also view statistics about volume bandwidth and I/O usage.

Related concepts

[Quality of Service](#) on page 69

A SolidFire storage cluster has the ability to provide Quality of Service (QoS) parameters on a per-volume basis. You can guarantee cluster performance measured in inputs and outputs per second (IOPS) using three configurable parameters that define QoS: Min IOPS, Max IOPS, and Burst IOPS.

Related tasks

[Creating a QoS policy](#) on page 72

You can create QoS policies and apply them when creating volumes.

[Editing a QoS policy](#) on page 72

You can change the name of an existing QoS policy or edit the values associated with the policy. Changing a QoS policy affects all volumes associated with the policy.

[Deleting a QoS policy](#) on page 73

You can delete a QoS policy if it is no longer needed. When you delete a QoS policy, all volumes associated with the policy maintain the QoS settings but become unassociated with a policy.

[Creating a volume](#) on page 73

You can create a volume and associate the volume with a given account. Every volume must be associated with an account. This association gives the account access to the volume through the iSCSI initiators using the CHAP credentials.

[Viewing individual volume performance details](#) on page 75

You can view performance statistics for individual volumes.

[Editing active volumes](#) on page 75

You can modify volume attributes such as QoS values, volume size, and the unit of measurement in which byte values are calculated. You can also modify account access for replication usage or to restrict access to the volume.

[Deleting a volume](#) on page 76

You can delete one or more volumes from an Element storage cluster.

[Restoring a deleted volume](#) on page 77

You can restore a volume in the system if it has been deleted but not yet purged. The system automatically purges a volume approximately eight hours after it has been deleted. If the system has purged the volume, you cannot restore it.

[Purging a volume](#) on page 77

When a volume is purged, it is permanently removed from the system. All data in the volume is lost.

[Cloning a volume](#) on page 77

You can create a clone of a single volume or multiple volumes to make a point-in-time copy of the data. When you clone a volume, the system creates a snapshot of the volume and then creates a copy of the data referenced by the snapshot. This is an asynchronous process, and the amount of

time the process requires depends on the size of the volume you are cloning and the current cluster load.

[*Assigning LUNs to Fibre Channel volumes*](#) on page 79

You can change the LUN assignment for a Fibre Channel volume in a volume access group. You can also make Fibre Channel volume LUN assignments when you create a volume access group.

[*Applying a QoS policy to volumes*](#) on page 79

You can bulk apply an existing QoS policy to one or more volumes.

[*Removing the QoS policy association of a volume*](#) on page 79

You can remove a QoS policy association from a volume by selecting custom QoS settings.

Quality of Service

A SolidFire storage cluster has the ability to provide Quality of Service (QoS) parameters on a per-volume basis. You can guarantee cluster performance measured in inputs and outputs per second (IOPS) using three configurable parameters that define QoS: Min IOPS, Max IOPS, and Burst IOPS.

Note: SolidFire Active IQ has a QoS recommendations page that provides advice on optimal configuration and set up of QoS settings.

IOPS parameters are defined in the following ways:

Minimum IOPS

The minimum number of sustained inputs and outputs per second (IOPS) that the storage cluster provides to a volume. The Min IOPS configured for a volume is the guaranteed level of performance for a volume. Performance does not drop below this level.

Maximum IOPS

The maximum number of sustained IOPS that the storage cluster provides to a volume. When cluster IOPS levels are critically high, this level of IOPS performance is not exceeded.

Burst IOPS

The maximum number of IOPS allowed in a short burst scenario. If a volume has been running below the Max IOPS, burst credits are accumulated. When performance levels become very high and are pushed to maximum levels, short bursts of IOPS are allowed on the volume.

Element software uses Burst IOPS when a cluster is running in a state of low cluster IOPS utilization.

A single volume can accrue Burst IOPS and use the credits to burst above their Max IOPS up to their Burst IOPS level for a set "burst period". A volume can burst for up to 60 seconds if the cluster has the capacity to accommodate the burst.

A volume accrues one second of burst credit (up to a maximum of 60 seconds) for every second that the volume runs below its Max IOPS limit.

Burst IOPS are limited in two ways:

- A volume can burst above its Max IOPS for a number of seconds equal to the number of burst credits that the volume has accrued.
- When a volume bursts above its Max IOPS setting, it is limited by its Burst IOPS setting. Therefore, the burst IOPS never exceeds the burst IOPS setting for the volume.

Effective Max Bandwidth

The maximum bandwidth is calculated by multiplying the number of IOPS (based on the QoS curve) by the IO size.

Example:

QoS parameter settings of 100 Min IOPS, 1000 Max IOPS, and 1500 Burst IOPs have the following effects on quality of performance:

- Workloads are able to reach and sustain a maximum of 1000 IOPS until the condition of workload contention for IOPS becomes apparent on the cluster. IOPS are then reduced incrementally until IOPS on all volumes are within the designated QoS ranges and contention for performance is relieved.
- Performance on all volumes is pushed toward the Min IOPS of 100. Levels do not drop below the Min IOPS setting but could remain higher than 100 IOPS when workload contention is relieved.
- Performance is never greater than 1000 IOPS, or less than 100 IOPS for a sustained period. Performance of 1500 IOPS (Burst IOPS) is allowed, but only for those volumes that have accrued burst credits by running below Max IOPS and only allowed for a short periods of time. Burst levels are never sustained.

QoS value limits

You can find information about the possible minimum and maximum values for Quality of Service (QoS).

			I/O size maximum value			
Parameter s	Minimum value	Default	4KB	8KB	16KB	262KB
Min IOPS	50	50	15,000	9,375*	5556*	385*
Max IOPS	100	15,000	200,000**	125,000	74,074	5128
Burst IOPS	100	15,000	200,000**	125,000	74,074	5128

*These estimations are approximate.

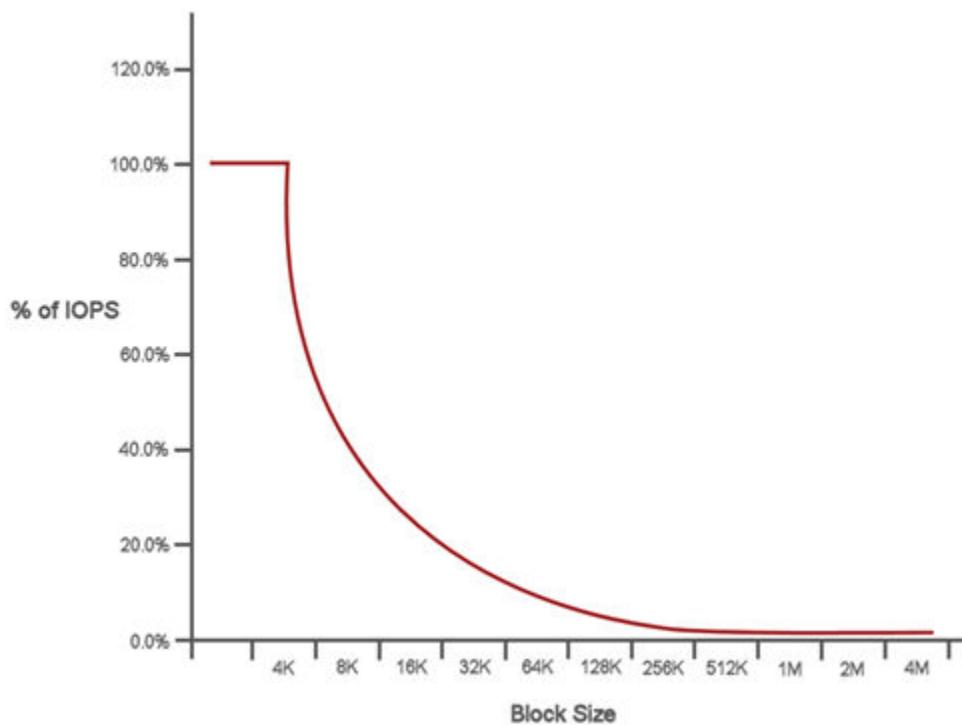
**Max IOPS and Burst IOPS can be set as high as 200,000; however, this setting is allowed only to effectively uncap the performance of a volume. Real-world maximum performance of a volume is limited by cluster usage and per-node performance.

QoS performance curve

The Quality of Service (QoS) performance curve shows the relationship between block size and the percentage of IOPS.

Block size and bandwidth have a direct impact on the number of IOPS that an application can obtain. Element software takes into account the block sizes it receives by normalizing block sizes to 4k. Based on workload, the system might increase block sizes. As block sizes increase, the system increases bandwidth to a level necessary to process the larger block sizes. As bandwidth increases the number of IOPS the system is able to attain decreases.

The QoS performance curve shows the relationship between increasing block sizes and the decreasing percentage of IOPS:



As an example, if block sizes are 4k, and bandwidth is 4000 KBps, the IOPS are 1000. If block sizes increase to 8k, bandwidth increases to 5000 KBps, and IOPS decrease to 625. By taking block size into account, the system ensures that lower priority workloads that use higher block sizes, such as backups and hypervisor activities, do not take too much of the performance needed by higher priority traffic using smaller block sizes.

QoS policies

A Quality of Service (QoS) policy enables you to create and save a standardized quality of service setting that can be applied to many volumes. You can create, edit, and delete QoS policies from the QoS Policies page on the Management tab.

Note: If you are using QoS policies, do not use custom QoS on a volume. Custom QoS will override and adjust QoS policy values for volume QoS settings.

[NetApp video: SolidFire Quality of Service Policies](#)

Related tasks

[Creating a QoS policy](#) on page 72

You can create QoS policies and apply them when creating volumes.

[Editing a QoS policy](#) on page 72

You can change the name of an existing QoS policy or edit the values associated with the policy. Changing a QoS policy affects all volumes associated with the policy.

[Deleting a QoS policy](#) on page 73

You can delete a QoS policy if it is no longer needed. When you delete a QoS policy, all volumes associated with the policy maintain the QoS settings but become unassociated with a policy.

Creating a QoS policy

You can create QoS policies and apply them when creating volumes.

Steps

1. Select **Management > QoS Policies**.
2. Click **Create QoS Policy**.
3. Enter the **Policy Name**.
4. Enter the **Min IOPS**, **Max IOPS**, and **Burst IOPS** values.
5. Click **Create QoS Policy**.

QoS policies details

You can view details of QoS policies from the Management tab.

ID

The system-generated ID for the QoS policy.

Name

The user-defined name for the QoS policy.

Min IOPS

The minimum number of IOPS guaranteed for the volume.

Max IOPS

The maximum number of IOPS allowed for the volume.

Burst IOPS

The maximum number of IOPS allowed over a short period of time for the volume. Default = 15,000.

Volumes

Shows the number of volumes using the policy. This number links to a table of volumes that have the policy applied.

Editing a QoS policy

You can change the name of an existing QoS policy or edit the values associated with the policy. Changing a QoS policy affects all volumes associated with the policy.

Steps

1. Select **Management > QoS Policies**.
2. Click the Actions icon for the QoS policy you want to edit.
3. In the resulting menu, select **Edit**.
4. In the **Edit QoS Policy** dialog box, modify the following properties as required:
 - Policy Name
 - Min IOPS
 - Max IOPS
 - Burst IOPS
5. Click **Save Changes**.

Deleting a QoS policy

You can delete a QoS policy if it is no longer needed. When you delete a QoS policy, all volumes associated with the policy maintain the QoS settings but become unassociated with a policy.

About this task

Note: If you are trying instead to disassociate a volume from a QoS policy, you can change the QoS settings for that volume to custom.

Steps

1. Select **Management > QoS Policies**.
2. Click the Actions icon for the QoS policy you want to delete.
3. In the resulting menu, select **Delete**.
4. Confirm the action.

Related tasks

[Removing the QoS policy association of a volume](#) on page 79

You can remove a QoS policy association from a volume by selecting custom QoS settings.

Creating a volume

You can create a volume and associate the volume with a given account. Every volume must be associated with an account. This association gives the account access to the volume through the iSCSI initiators using the CHAP credentials.

About this task

You can specify QoS settings for a volume during creation.

Steps

1. Select **Management > Volumes**.
2. Click **Create Volume**.
3. In the **Create a New Volume** dialog box, enter the **Volume Name**.
4. Enter the total size of the volume.

Note: The default volume size selection is in GB. You can create volumes using sizes measured in GB or GiB:

- 1GB = 1 000 000 000 bytes
- 1GiB = 1 073 741 824 bytes

5. Select a **Block Size** for the volume.

6. Click the **Account** drop-down list and select the account that should have access to the volume.

If an account does not exist, click the **Create Account** link, enter a new account name, and click **Create**. The account is created and associated with the new volume.

Note: If there are more than 50 accounts, the list does not appear. Begin typing and the auto-complete function displays possible values for you to choose.

7. To set the **Quality of Service**, do one of the following:

- a. Under **Policy**, you can select an existing QoS policy, if available.
- b. Under **Custom Settings**, set customized minimum, maximum, and burst values for IOPS or use the default QoS values.

Volumes that have a Max or Burst IOPS value greater than 20,000 IOPS might require high queue depth or multiple sessions to achieve this level of IOPS on a single volume.

8. Click **Create Volume**.

Volume details

The Volumes page on the Management tab provides information about the active volumes such as name, account, access groups associated with the volume, and size of the volume.

ID

The system-generated ID for the volume.

Name

The name given to the volume when it was created.

Account

The name of the account assigned to the volume.

Access Groups

The name of the volume access group or groups to which the volume belongs.

Access

The type of access assigned to the volume when it was created. Possible values:

- Read / Write: All reads and writes are accepted.
- Read Only: All read activity allowed; no writes allowed.
- Locked: Only Administrator access allowed.
- ReplicationTarget: Designated as a target volume in a replicated volume pair.

Used

The percentage of used space in the volume.

Size

The total size (in GB) of the volume.

Snapshots

The number of snapshots created for the volume.

QoS Policy

The name and link to the user-defined QoS policy.

Min IOPS

The minimum number of IOPS guaranteed for the volume.

Max IOPS

The maximum number of IOPS allowed for the volume.

Burst IOPS

The maximum number of IOPS allowed over a short period of time for the volume. Default = 15,000.

Attributes

Attributes that have been assigned to the volume as a key/value pair through an API method.

512e

Indication of whether 512e is enabled on a volume. Possible values:

- Yes
- No

Created On

The date and time that the volume was created.

Viewing individual volume performance details

You can view performance statistics for individual volumes.

Steps

1. Select **Reporting > Volume Performance**.
2. In the volume list, click the Actions icon for a volume.
3. Click **View Details**.
A tray appears at the bottom of the page containing general information about the volume.
4. To see more detailed information about the volume, click **See More Details**.
The system displays detailed information as well as performance graphs for the volume.

Editing active volumes

You can modify volume attributes such as QoS values, volume size, and the unit of measurement in which byte values are calculated. You can also modify account access for replication usage or to restrict access to the volume.

About this task

You can resize a volume when there is sufficient space on the cluster under the following conditions:

- Normal operating conditions.
- Volume errors or failures are being reported.
- The volume is being cloned.
- The volume is being resynced.

Steps

1. Select **Management > Volumes**.
2. In the **Active** window, click the Actions icon for the volume you want to edit.
3. Click **Edit**.
4. Optional: Change the total size of the volume.

Note:

- You can increase, but not decrease, the size of the volume. You can only resize one volume in a single resizing operation. Garbage collection operations and software upgrades do not interrupt the resizing operation.
- If you are adjusting volume size for replication, you should first increase the size of the volume assigned as the replication target. Then you can resize the source volume. The target volume can be greater or equal in size to the source volume, but it cannot be smaller.

Note: The default volume size selection is in GB. You can create volumes using sizes measured in GB or GiB:

- 1GB = 1 000 000 000 bytes
- 1GiB = 1 073 741 824 bytes

5. Optional: Select a different account access level of one of the following:
 - Read Only

- Read/Write
- Locked
- Replication Target

6. Optional: Select the account that should have access to the volume.
If the account does not exist, click the **Create Account** link, enter a new account name, and click **Create**. The account is created and associated with the volume.

Note: If there are more than 50 accounts, the list does not appear. Begin typing and the auto-complete function displays possible values for you to choose.

7. Optional: To change the selection in **Quality of Service**, do one of the following:

- a. Under **Policy**, you can select an existing QoS policy, if available.
- b. Under **Custom Settings**, set customized minimum, maximum, and burst values for IOPS or use the default QoS values.

Note: If you are using QoS policies on a volume, you can set custom QoS to remove the QoS policy affiliation with the volume. Custom QoS will override and adjust QoS policy values for volume QoS settings.

Tip: When you change IOPS values, you should increment in tens or hundreds. Input values require valid whole numbers.

Tip: Configure volumes with an extremely high burst value. This allows the system to process occasional large block sequential workloads more quickly, while still constraining the sustained IOPS for a volume.

8. Click **Save Changes**.

Deleting a volume

You can delete one or more volumes from an Element storage cluster.

About this task

The system does not immediately purge a deleted volume; the volume remains available for approximately eight hours. If you restore a volume before the system purges it, the volume comes back online and iSCSI connections are restored.

If a volume used to create a snapshot is deleted, its associated snapshots become inactive. When the deleted source volumes are purged, the associated inactive snapshots are also removed from the system.



Attention: Persistent volumes that are associated with management services are created and assigned to a new account during installation or upgrade. If you are using persistent volumes, do not modify or delete the volumes or their associated account.

Steps

1. Select **Management > Volumes**.
2. To delete a single volume, perform the following steps:
 - a. Click the Actions icon for the volume you want to delete.
 - b. In the resulting menu, click **Delete**.
 - c. Confirm the action.

The system moves the volume to the **Deleted** area on the **Volumes** page.
3. To delete multiple volumes, perform the following steps:

- a. In the list of volumes, check the box next to any volumes you want to delete.
- b. Click **Bulk Actions**.
- c. In the resulting menu, click **Delete**.
- d. Confirm the action.

The system moves the volumes to the **Deleted** area on the **Volumes** page.

Restoring a deleted volume

You can restore a volume in the system if it has been deleted but not yet purged. The system automatically purges a volume approximately eight hours after it has been deleted. If the system has purged the volume, you cannot restore it.

Steps

1. Select **Management > Volumes**.
2. Click the **Deleted** tab to view the list of deleted volumes.
3. Click the Actions icon for the volume you want to restore.
4. In the resulting menu, click **Restore**.
5. Confirm the action.

The volume is placed in the **Active** volumes list and iSCSI connections to the volume are restored.

Purging a volume

When a volume is purged, it is permanently removed from the system. All data in the volume is lost.

About this task

The system automatically purges deleted volumes eight hours after deletion. However, if you want to purge a volume before the scheduled time, you can do so.

Steps

1. Select **Management > Volumes**.
2. Click the **Deleted** button.
3. Perform the steps to purge a single volume or multiple volumes.

Option	Steps
Purge a single volume	<ol style="list-style-type: none">a. Click the Actions icon for the volume you want to purge.b. Click Purge.c. Confirm the action.
Purge multiple volumes	<ol style="list-style-type: none">a. Select the volumes you want to purge.b. Click Bulk Actions.c. In the resulting menu, select Purge.d. Confirm the action.

Cloning a volume

You can create a clone of a single volume or multiple volumes to make a point-in-time copy of the data. When you clone a volume, the system creates a snapshot of the volume and then creates a copy of the data referenced by the snapshot. This is an asynchronous process, and the amount of time the process requires depends on the size of the volume you are cloning and the current cluster load.

About this task

The cluster supports up to two running clone requests per volume at a time and up to eight active volume clone operations at a time. Requests beyond these limits are queued for later processing.

Note: Operating systems differ in how they treat cloned volumes. VMware ESXi will treat a cloned volume as a volume copy or snapshot volume. The volume will be an available device to use to create a new datastore. For more information on mounting clone volumes and handling snapshot LUNs, see VMware documentation on [mounting a VMFS datastore copy](#) and [managing duplicate VMFS datastores](#).



Attention: Before you truncate a cloned volume by cloning to a smaller size, ensure that you prepare the partitions so that they fit into the smaller volume.

Steps

1. Select **Management > Volumes**.

2. To clone a single volume, perform the following steps:

- a. In the list of volumes on the **Active** page, click the Actions icon for the volume you want to clone.
- b. In the resulting menu, click **Clone**.
- c. In the **Clone Volume** window, enter a volume name for the newly cloned volume.
- d. Select a size and measurement for the volume using the **Volume Size** spin box and list.

Note: The default volume size selection is in GB. You can create volumes using sizes measured in GB or GiB:

- 1GB = 1 000 000 000 bytes
- 1GiB = 1 073 741 824 bytes

- e. Select the type of access for the newly cloned volume.
- f. Select an account to associate with the newly cloned volume from the **Account** list.

Note: You can create an account during this step if you click the **Create Account** link, enter an account name, and click **Create**. The system automatically adds the account to the **Account** list after you create it.

3. To clone multiple volumes, perform the following steps:

- a. In the list of volumes on the **Active** page, check the box next to any volumes you want to clone.
- b. Click **Bulk Actions**.
- c. In the resulting menu, select **Clone**.
- d. In the **Clone Multiple Volumes** dialog box, enter a prefix for the cloned volumes in the **New Volume Name Prefix** field.
- e. Select an account to associate with the cloned volumes from the **Account** list.
- f. Select the type of access for the cloned volumes.

4. Click **Start Cloning**.

Note: Increasing the volume size of a clone results in a new volume with additional free space at the end of the volume. Depending on how you use the volume, you might need to extend partitions or create new partitions in the free space to make use of it.

Assigning LUNs to Fibre Channel volumes

You can change the LUN assignment for a Fibre Channel volume in a volume access group. You can also make Fibre Channel volume LUN assignments when you create a volume access group.

About this task

Assigning new Fibre Channel LUNs is an advanced function and could have unknown consequences on the connecting host. For example, the new LUN ID might not be automatically discovered on the host, and the host might require a rescan to discover the new LUN ID.

Steps

1. Select **Management > Access Groups**.
2. Click the Actions icon for the access group you want to edit.
3. In the resulting menu, select **Edit**.
4. Under **Assign LUN IDs** in the **Edit Volume Access Group** dialog box, click the arrow on the **LUN Assignments** list.
5. For each volume in the list that you want to assign a LUN to, enter a new value in the corresponding **LUN** field.
6. Click **Save Changes**.

Applying a QoS policy to volumes

You can bulk apply an existing QoS policy to one or more volumes.

Before you begin

The QoS policy you want to bulk apply exists.

Steps

1. Select **Management > Volumes**.
2. In the list of volumes, check the box next to any volumes you want to apply the QoS policy to.
3. Click **Bulk Actions**.
4. In the resulting menu, click **Apply QoS Policy**.
5. Select the QoS policy from the drop-down list.
6. Click **Apply**.

Related concepts

[QoS policies](#) on page 71

A Quality of Service (QoS) policy enables you to create and save a standardized quality of service setting that can be applied to many volumes. You can create, edit, and delete QoS policies from the QoS Policies page on the Management tab.

Removing the QoS policy association of a volume

You can remove a QoS policy association from a volume by selecting custom QoS settings.

Before you begin

The volume you want to modify is associated with a QoS policy.

Steps

1. Select **Management > Volumes**.
2. Click the Actions icon for a volume that contains a QoS policy you want to modify.
3. Click **Edit**.
4. In the resulting menu under **Quality of Service**, click **Custom Settings**.

5. Modify **Min IOPS**, **Max IOPS**, and **Burst IOPS**, or keep the default settings.
6. Click **Save Changes**.

Related tasks

[Deleting a QoS policy](#) on page 73

You can delete a QoS policy if it is no longer needed. When you delete a QoS policy, all volumes associated with the policy maintain the QoS settings but become unassociated with a policy.

Working with virtual volumes

You can view information and perform tasks for virtual volumes and their associated storage containers, protocol endpoints, bindings, and hosts using the Element UI.

The NetApp Element software storage system ships with the Virtual Volumes (VVols) feature disabled. You must perform a one-time task of manually enabling vSphere VVol functionality through the Element UI.

After you enable the VVol functionality, a VVols tab appears in the user interface that offers VVols-related monitoring and limited management options. Additionally, a storage-side software component known as the VASA Provider acts as a storage awareness service for vSphere. Most VVols commands, such as VVol creation, cloning, and editing, are initiated by a vCenter Server or ESXi host and translated by the VASA Provider to Element APIs for the Element software storage system. Commands to create, delete, and manage storage containers and delete virtual volumes can be initiated using the Element UI.

The majority of configurations necessary for using Virtual Volumes functionality with Element software storage systems are made in vSphere. See the *VMware vSphere Virtual Volumes for SolidFire Storage Configuration Guide* to register the VASA Provider in vCenter, create and manage VVol datastores, and manage storage based on policies.

Note: VASA support for multiple vCenters is available as an upgrade patch if you have already registered a VASA provider with your vCenter. To install, follow the directions in the VASA39 manifest and download the .tar.gz file from the [NetApp Software Downloads](#) site. The NetApp Element VASA provider uses a NetApp certificate. With this patch, the certificate is used unmodified by vCenter to support multiple vCenters for VASA and VVols use. Do not modify the certificate. Custom SSL certificates are not supported by VASA.

Related concepts

[Protocol endpoints](#) on page 87

Protocol endpoints are access points used by a host to address storage on a cluster running NetApp Element software. Protocol endpoints cannot be deleted or modified by a user, are not associated with an account, and cannot be added to a volume access group.

[Bindings](#) on page 88

To perform I/O operations with a virtual volume, an ESXi host must first bind the virtual volume.

Related tasks

[Enabling virtual volumes](#) on page 81

You must manually enable vSphere Virtual Volumes (VVols) functionality through the NetApp Element software. The Element software system comes with VVols functionality disabled by default, and it is not automatically enabled as part of a new installation or upgrade. Enabling the VVols feature is a one-time configuration task.

[Deleting a virtual volume](#) on page 84

Although virtual volumes should always be deleted from the VMware Management Layer, the functionality for you to delete virtual volumes is enabled from the Element UI. You should only

delete a virtual volume from the Element UI when absolutely necessary, such as when vSphere fails to clean up virtual volumes on SolidFire storage.

[Creating a storage container](#) on page 85

You can create storage containers in the Element UI and discover them in vCenter. You must create at least one storage container to begin provisioning VVol-backed virtual machines.

[Editing a storage container](#) on page 86

You can modify storage container CHAP authentication in the Element UI.

[Deleting a storage container](#) on page 87

You can delete storage containers from the Element UI.

Related reference

[Host details](#) on page 88

The Hosts page on the VVols tab provides information about VMware ESXi hosts that host virtual volumes.

Enabling virtual volumes

You must manually enable vSphere Virtual Volumes (VVols) functionality through the NetApp Element software. The Element software system comes with VVols functionality disabled by default, and it is not automatically enabled as part of a new installation or upgrade. Enabling the VVols feature is a one-time configuration task.

Before you begin

- The cluster must be running Element 9.0 or later.
- The cluster must be connected to an ESXi 6.0 or later environment that is compatible with VVols.
- If you are using Element 11.3 or later, the cluster must be connected to an ESXi 6.0 update 3 or later environment.

About this task



Attention: Enabling vSphere Virtual Volumes functionality permanently changes the Element software configuration. You should only enable VVols functionality if your cluster is connected to a VMware ESXi VVols-compatible environment. You can disable the VVols feature and restore the default settings only by returning the cluster to the factory image, which deletes all data on the system.

Steps

1. Select **Clusters > Settings**.
2. Find the cluster-specific settings for Virtual Volumes.
3. Click **Enable Virtual Volumes**.
4. Click **Yes** to confirm the Virtual Volumes configuration change.
The **VVols** tab appears in the Element UI.

Note: When VVols functionality is enabled, the SolidFire cluster starts the VASA Provider, opens port 8444 for VASA traffic, and creates protocol endpoints that can be discovered by vCenter and all ESXi hosts.

5. Copy the VASA Provider URL from the Virtual Volumes (VVols) settings in **Clusters > Settings**. You will use this URL to register the VASA Provider in vCenter.
6. Create a storage container in **VVols > Storage Containers**.

Note: You must create at least one storage container so that VMs can be provisioned to a VVol datastore.

7. Select **VVols > Protocol Endpoints**.
8. Verify that a protocol endpoint has been created for each node in the cluster.

Note: Additional configuration tasks are required in vSphere. See the *VMware vSphere Virtual Volumes for SolidFire Storage Configuration Guide* to register the VASA Provider in vCenter, create and manage VVol datastores, and manage storage based on policies.

Related information

[VMware vSphere Virtual Volumes for SolidFire Storage Configuration Guide](#)

Viewing virtual volume details

You can review virtual volume information for all active virtual volumes on the cluster in the Element UI. You can also view performance activity for each virtual volume, including input, output, throughput, latency, queue depth, and volume information.

Before you begin

- You have enabled VVols functionality in the Element UI for the cluster.
- You have created an associated storage container.
- You have configured your vSphere cluster to use Element software VVols functionality.
- You have created at least one VM in vSphere.

Steps

1. Click **VVols > Virtual Volumes**.
The information for all active virtual volumes is displayed.
2. Click the Actions icon for the virtual volume you want to review.
3. In the resulting menu, select **View Details**.

Virtual volume details

The Virtual Volumes page of the VVols tab provides information about each active virtual volume on the cluster, such as volume ID, snapshot ID, parent virtual volume ID, and virtual volume ID.

Volume ID

The ID of the underlying volume.

Snapshot ID

The ID of the underlying volume snapshot. The value is 0 if the virtual volume does not represent a SolidFire snapshot.

Parent Virtual Volume ID

The virtual volume ID of the parent virtual volume. If the ID is all zeros, the virtual volume is independent with no link to a parent.

Virtual Volume ID

The UUID of the virtual volume.

Name

The name assigned to the virtual volume.

Storage Container

The storage container that owns the virtual volume.

Guest OS Type

Operating system associated with the virtual volume.

Virtual Volume Type

The virtual volume type: Config, Data, Memory, Swap, or Other.

Access

The read-write permissions assigned to the virtual volume.

Size

The size of the virtual volume in GB or GiB.

Snapshots

The number of associated snapshots. Click the number to link to snapshot details.

Min IOPS

The minimum IOPS QoS setting of the virtual volume.

Max IOPS

The maximum IOPS QoS setting of the virtual volume.

Burst IOPS

The maximum burst QoS setting of the virtual volume.

VMW_VmID

Information in fields prefaced with "VMW_" are defined by VMware.

Create Time

The time the virtual volume creation task was completed.

Individual virtual volume details

The Virtual Volumes page on the VVols tab provides the following virtual volume information when you select an individual virtual volume and view its details.

VMW_XXX

Information in fields prefaced with "VMW_" are defined by VMware.

Parent Virtual Volume ID

The virtual volume ID of the parent virtual volume. If the ID is all zeros, the virtual volume is independent with no link to a parent.

Virtual Volume ID

The UUID of the virtual volume.

Virtual Volume Type

The virtual volume type: Config, Data, Memory, Swap, or Other.

Volume ID

The ID of the underlying volume.

Access

The read-write permissions assigned to the virtual volume.

Account Name

Name of the account containing the volume.

Access Groups

Associated volume access groups.

Total Volume Size

Total provisioned capacity in bytes.

Non-Zero Blocks

Total number of 4KiB blocks with data after the last garbage collection operation has completed.

Zero Blocks

Total number of 4KiB blocks without data after the last round of garbage collection operation has completed.

Snapshots

The number of associated snapshots. Click the number to link to snapshot details.

Min IOPS

The minimum IOPS QoS setting of the virtual volume.

Max IOPS

The maximum IOPS QoS setting of the virtual volume.

Burst IOPS

The maximum burst QoS setting of the virtual volume.

Enable 512

Because virtual volumes always use 512-byte block size emulation, the value is always yes.

Volumes Paired

Indicates if a volume is paired.

Create Time

The time the virtual volume creation task was completed.

Blocks Size

Size of the blocks on the volume.

Unaligned Writes

For 512e volumes, the number of write operations that were not on a 4k sector boundary.

High numbers of unaligned writes might indicate improper partition alignment.

Unaligned Reads

For 512e volumes, the number of read operations that were not on a 4k sector boundary.

High numbers of unaligned reads might indicate improper partition alignment.

scsiEUIDeviceID

Globally unique SCSI device identifier for the volume in EUI-64 based 16-byte format.

scsiNAADeviceID

Globally unique SCSI device identifier for the volume in NAA IEEE Registered Extended format.

Attributes

List of name-value pairs in JSON object format.

Deleting a virtual volume

Although virtual volumes should always be deleted from the VMware Management Layer, the functionality for you to delete virtual volumes is enabled from the Element UI. You should only delete a virtual volume from the Element UI when absolutely necessary, such as when vSphere fails to clean up virtual volumes on SolidFire storage.

Steps

1. Select **VVols > Virtual Volumes**.

2. Click the Actions icon for the virtual volume you want to delete.
3. In the resulting menu, select **Delete**.



Attention: You should delete a virtual volume from the VMware Management Layer to ensure that the virtual volume is properly unbound before deletion. You should only delete a virtual volume from the Element UI when absolutely necessary, such as when vSphere fails to clean up virtual volumes on SolidFire storage. If you delete a virtual volume from the Element UI, the volume will be purged immediately.

4. Confirm the action.
5. Refresh the list of virtual volumes to confirm that the virtual volume has been removed.
6. Optional: Select **Reporting > Event Log** to confirm that the purge has been successful.

Storage containers

A storage container is a vSphere datastore representation created on a cluster running Element software.

Storage containers are created and tied to NetApp Element accounts. A storage container created on Element storage appears as a vSphere datastore in vCenter and ESXi. Storage containers do not allocate any space on Element storage. They are simply used to logically associate virtual volumes.

A maximum of four storage containers per cluster is supported. A minimum of one storage container is required to enable VVols functionality.

Creating a storage container

You can create storage containers in the Element UI and discover them in vCenter. You must create at least one storage container to begin provisioning VVol-backed virtual machines.

Before you begin

You have enabled VVols functionality in the Element UI for the cluster.

Steps

1. Select **VVols > Storage Containers**.
2. Click the **Create Storage Containers** button.
3. Enter storage container information in the **Create a New Storage Container** dialog box:
 - a. Enter a name for the storage container.
 - b. Configure initiator and target secrets for CHAP.

Tip: Leave the CHAP Settings fields blank to automatically generate secrets.

- c. Click the **Create Storage Container** button.
4. Verify that the new storage container appears in the list in the **Storage Containers** sub-tab.

Note: Because a NetApp Element account ID is created automatically and assigned to the storage container, it is not necessary to manually create an account.

Storage container details

On the Storage Containers page of the VVols tab, you can view information for all active storage containers on the cluster.

Account ID

The ID of the NetApp Element account associated with the storage container.

Name

The name of the storage container.

Status

The status of the storage container. Possible values:

- **Active:** The storage container is in use.
- **Locked:** The storage container is locked.

PE Type

The protocol endpoint type (SCSI is the only available protocol for Element software).

Storage Container ID

The UUID of the virtual volume storage container.

Active Virtual Volumes

The number of active virtual volumes associated with the storage container.

Individual storage container details

You can view the storage container information for an individual storage container by selecting it from the Storage Containers page on the VVols tab.

Account ID

The ID of the NetApp Element account associated with the storage container.

Name

The name of the storage container.

Status

The status of the storage container. Possible values:

- **Active:** The storage container is in use.
- **Locked:** The storage container is locked.

Chap Initiator Secret

The unique CHAP secret for the initiator.

Chap Target Secret

The unique CHAP secret for the target.

Storage Container ID

The UUID of the virtual volume storage container.

Protocol Endpoint Type

Indicates the protocol endpoint type (SCSI is the only available protocol).

Editing a storage container

You can modify storage container CHAP authentication in the Element UI.

Steps

1. Select **VVols > Storage Containers**.
2. Click the Actions icon for the storage container you want to edit.

3. In the resulting menu, select **Edit**.
4. Under CHAP Settings, edit the Initiator Secret and Target Secret credentials used for authentication.

Tip: If you do not change the CHAP Settings credentials, they remain the same. If you make the credentials fields blank, the system automatically generates new secrets.

5. Click **Save Changes**.

Deleting a storage container

You can delete storage containers from the Element UI.

Before you begin

All virtual machines have been removed from the VVol datastore.

Steps

1. Select **VVols > Storage Containers**.
2. Click the Actions icon for the storage container you want to delete.
3. In the resulting menu, select **Delete**.
4. Confirm the action.
5. Refresh the list of storage containers in the **Storage Containers** sub-tab to confirm that the storage container has been removed.

Protocol endpoints

Protocol endpoints are access points used by a host to address storage on a cluster running NetApp Element software. Protocol endpoints cannot be deleted or modified by a user, are not associated with an account, and cannot be added to a volume access group.

A cluster running Element software automatically creates one protocol endpoint per storage node in the cluster. For example, a six-node storage cluster has six protocol endpoints that are mapped to each ESXi host. Protocol endpoints are dynamically managed by Element software and are created, moved, or removed as needed without any intervention. Protocol endpoints are the target for multi-pathing and act as an I/O proxy for subsidiary LUNs. Each protocol endpoint consumes an available SCSI address, just like a standard iSCSI target. Protocol endpoints appear as a single-block (512-byte) storage device in the vSphere client, but this storage device is not available to be formatted or used as storage.

iSCSI is the only supported protocol. Fibre Channel protocol is not supported.

Protocol endpoints details

The Protocol Endpoints page on the VVols tab provides protocol endpoint information.

Primary Provider ID

The ID of the primary protocol endpoint provider.

Secondary Provider ID

The ID of the secondary protocol endpoint provider.

Protocol Endpoint ID

The UUID of the protocol endpoint.

Protocol Endpoint State

The status of the protocol endpoint. Possible values are as follows:

- **Active:** The protocol endpoint is in use.
- **Start:** The protocol endpoint is starting.

- **Failover:** The protocol endpoint has failed over.
- **Reserved:** The protocol endpoint is reserved.

Provider Type

The type of the protocol endpoint's provider. Possible values are as follows:

- **Primary**
- **Secondary**

SCSI NAA Device ID

The globally unique SCSI device identifier for the protocol endpoint in NAA IEEE Registered Extended Format.

Bindings

To perform I/O operations with a virtual volume, an ESXi host must first bind the virtual volume.

The SolidFire cluster chooses an optimal protocol endpoint, creates a binding that associates the ESXi host and virtual volume with the protocol endpoint, and returns the binding to the ESXi host. After it is bound, the ESXi host can perform I/O operations with the bound virtual volume.

Bindings details

The Bindings page on the VVols tab provides binding information about each virtual volume.

The following information is displayed:

Host ID

The UUID for the ESXi host that hosts virtual volumes and is known to the cluster.

Protocol Endpoint ID

Protocol endpoint IDs that correspond to each node in the SolidFire cluster.

Protocol Endpoint in Band ID

The SCSI NAA device ID of the protocol endpoint.

Protocol Endpoint Type

The protocol endpoint type.

VVol Binding ID

The binding UUID of the virtual volume.

VVol ID

The universally unique identifier (UUID) of the virtual volume.

VVol Secondary ID

The secondary ID of the virtual volume that is a SCSI second level LUN ID.

Host details

The Hosts page on the VVols tab provides information about VMware ESXi hosts that host virtual volumes.

The following information is displayed:

Host ID

The UUID for the ESXi host that hosts virtual volumes and is known to the cluster.

Host Address

The IP address or DNS name for the ESXi host.

Bindings

Binding IDs for all virtual volumes bound by the ESXi host.

ESX Cluster ID

The vSphere host cluster ID or vCenter GUID.

Initiator IQNs

Initiator IQNs for the virtual volume host.

SolidFire Protocol Endpoint IDs

The protocol endpoints that are currently visible to the ESXi host.

Working with volume access groups and initiators

You can use iSCSI initiators or Fibre Channel initiators to access the volumes defined within volume access groups.

You can create access groups by mapping iSCSI initiator IQNs or Fibre Channel WWPNs in a collection of volumes. Each IQN that you add to an access group can access each volume in the group without requiring CHAP authentication.

There are two types of CHAP authentication methods:

- Account-level CHAP authentication: You can assign CHAP authentication for the account.
- Initiator-level CHAP authentication: You can assign unique CHAP target and secrets for specific initiators without being bound to single CHAP across a single account. This CHAP level authentication supersedes account level credentials.

Optionally, with per-initiator CHAP, you can enforce initiator authorization and per-initiator CHAP authentication. These options can be defined on a per-initiator basis and an access group can contain a mix of initiators with different options.

Each WWPN that you add to an access group enables Fibre Channel network access to the volumes in the access group.

Note: Volume access groups have the following limits:

- A maximum of 64 IQNs or WWPNs are allowed in an access group.
- An access group can be made up of a maximum of 2000 volumes.
- An IQN or WWPN can belong to only one access group.
- A single volume can belong to a maximum of four access groups.

Related tasks

[Creating a volume access group](#) on page 90

You can create volume access groups by mapping initiators to a collection of volumes for secured access. You can then grant access to the volumes in the group with an account CHAP initiator secret and target secret.

[Adding volumes to an access group](#) on page 92

You can add volumes to a volume access group. Each volume can belong to more than one volume access group; you can see the groups that each volume belongs to on the **Active** volumes page.

[Removing volumes from an access group](#) on page 92

When you remove a volume from an access group, the group no longer has access to that volume.

[Creating an initiator](#) on page 93

You can create iSCSI or Fibre Channel initiators and optionally assign them aliases.

[Editing an initiator](#) on page 93

You can change the alias of an existing initiator or add an alias if one does not already exist.

[Adding a single initiator to a volume access group](#) on page 94

You can add an initiator to an existing volume access group.

[Adding multiple initiators to a volume access group](#) on page 95

You can add multiple initiators to an existing volume access group to allow access to volumes in the volume access group with or without requiring CHAP authentication..

[Removing initiators from an access group](#) on page 95

When you remove an initiator from an access group, it can no longer access the volumes in that volume access group. Normal account access to the volume is not disrupted.

[Deleting an access group](#) on page 96

You can delete an access group when it is no longer needed. You do not need to delete Initiator IDs and Volume IDs from the volume access group before deleting the group. After you delete the access group, group access to the volumes is discontinued.

[Deleting an initiator](#) on page 96

You can delete an initiator after it is no longer needed. When you delete an initiator, the system removes it from any associated volume access group. Any connections using the initiator remain valid until the connection is reset.

Creating a volume access group

You can create volume access groups by mapping initiators to a collection of volumes for secured access. You can then grant access to the volumes in the group with an account CHAP initiator secret and target secret.

About this task

If you use initiator-based CHAP, you can add CHAP credentials for a single initiator in a volume access group, providing more security. This enables you to apply this option for volume access groups that already exist.

Steps

1. Click **Management > Access Groups**.
2. Click **Create Access Group**.
3. Enter a name for the volume access group in the **Name** field.
4. Add an initiator to the volume access group in one of the following ways:

Option	Description
Adding a Fibre Channel initiator	<ol style="list-style-type: none">Under Add Initiators, select an existing Fibre Channel initiator from the Unbound Fibre Channel Initiators list.Click Add FC Initiator. <p>Note: You can create an initiator during this step if you click the Create Initiator link, enter an initiator name, and click Create. The system automatically adds the initiator to the Initiators list after you create it. A sample of the format is as follows:</p>

5f:47:ac:c0:5c:74:d4:02

Option	Description
Adding an iSCSI initiator	<p>Under Add Initiators, select an existing initiator from the Initiators list.</p> <p>Note: You can create an initiator during this step if you click the Create Initiator link, enter an initiator name, and click Create. The system automatically adds the initiator to the Initiators list after you create it. A sample of the format is as follows:</p> <pre>iqn.2010-01.com.solidfire:c2r9.fc0.2100000e1e09bb8b</pre> <p>Tip: You can find the initiator IQN for each volume by selecting View Details in the Actions menu for the volume on the Management > Volumes > Active list.</p> <p>When you modify an initiator, you can toggle the <code>requiredCHAP</code> attribute to <code>True</code>, which enables you to set the target initiator secret. For details, see API information about the <code>ModifyInitiator</code> API method.</p>

[Managing storage with the Element API](#)

5. Optional: Add more initiators as needed.
6. Under Add Volumes, select a volume from the **Volumes** list.
The volume appears in the **Attached Volumes** list.
7. Optional: Add more volumes as needed.
8. Click **Create Access Group**.

Related tasks

[Adding volumes to an access group](#) on page 92

You can add volumes to a volume access group. Each volume can belong to more than one volume access group; you can see the groups that each volume belongs to on the **Active** volumes page.

Volume access group details

The Access Groups page on the Management tab provides information about volume access groups.

The following information is displayed:

ID

The system-generated ID for the access group.

Name

The name given to the access group when it was created.

Active Volumes

The number of active volumes in the access group.

Compression

The compression efficiency score for the access group.

Deduplication

The deduplication efficiency score for the access group.

Thin Provisioning

The thin provisioning efficiency score for the access group.

Overall Efficiency

The overall efficiency score for the access group.

Initiators

The number of initiators connected to the access group.

Viewing individual access group details

You can view details for an individual access group, such as attached volumes and initiators, in a graphical format.

Steps

1. Click **Management > Access Groups**.
2. Click the Actions icon for an access group.
3. Click **View Details**.

Adding volumes to an access group

You can add volumes to a volume access group. Each volume can belong to more than one volume access group; you can see the groups that each volume belongs to on the **Active** volumes page.

About this task

You can also use this procedure to add volumes to a Fibre Channel volume access group.

Steps

1. Click **Management > Access Groups**.
2. Click the Actions icon for the access group you want to add volumes to.
3. Click the **Edit** button.
4. Under Add Volumes, select a volume from the **Volumes** list.
You can add more volumes by repeating this step.
5. Click **Save Changes**.

Removing volumes from an access group

When you remove a volume from an access group, the group no longer has access to that volume.

About this task

Modifying CHAP settings in an account or removing initiators or volumes from an access group can cause initiators to lose access to volumes unexpectedly. To verify that volume access will not be lost unexpectedly, always logout iSCSI sessions that will be affected by an account or access group change, and verify that initiators can reconnect to volumes after any changes to initiator settings and cluster settings have been completed.

Steps

1. Click **Management > Access Groups**.
2. Click the Actions icon for the access group you want to remove volumes from.
3. Click **Edit**.
4. Under Add Volumes in the **Edit Volume Access Group** dialog box, click the arrow on the **Attached Volumes** list.
5. Select the volume you want to remove from the list and click the **x** icon to remove the volume from the list.
You can remove more volumes by repeating this step.
6. Click **Save Changes**.

Creating an initiator

You can create iSCSI or Fibre Channel initiators and optionally assign them aliases.

About this task

You can also assign initiator-based CHAP attributes by using an API call. To add a CHAP account name and credentials per initiator, you must use the `CreateInitiator` API call to remove and add CHAP access and attributes. For details, see the API reference information.

[Managing storage with the Element API](#)

Steps

1. Click **Management > Initiators**.
2. Click **Create Initiator**.
3. Perform the steps to create a single initiator or multiple initiators:

Option	Steps
Create a single initiator	<ol style="list-style-type: none">a. Click Create a Single Initiator.b. Enter the IQN or WWPN for the initiator in the IQN/WWPN field.c. Enter a friendly name for the initiator in the Alias field.d. Click Create Initiator.
Create multiple initiators	<ol style="list-style-type: none">a. Click Bulk Create Initiators.b. Enter a list of IQNs or WWPNs in the text box.c. Click Add Initiators.d. Choose an initiator from the resulting list and click the corresponding Add icon in the Alias column to add an alias for the initiator.e. Click the check mark to confirm the new alias.f. Click Create Initiators.

Editing an initiator

You can change the alias of an existing initiator or add an alias if one does not already exist.

About this task

To add a CHAP account name and credentials per initiator, you must use the `ModifyInitiator` API call to remove and add CHAP access and attributes. See API information for more details.

[Managing storage with the Element API](#)

Steps

1. Click **Management > Initiators**.
2. Click the Actions icon for the initiator you want to edit.
3. Click **Edit**.
4. Enter a new alias for the initiator in the **Alias** field.
5. Click **Save Changes**.

Adding a single initiator to a volume access group

You can add an initiator to an existing volume access group.

About this task

When you add an initiator to a volume access group, the initiator has access to all volumes in that volume access group.

Tip: You can find the initiator for each volume by clicking the Actions icon and then selecting **View Details** for the volume in the active volumes list.

If you use initiator-based CHAP, you can add CHAP credentials for a single initiator in a volume access group, providing more security. This enables you to apply this option for volume access groups that already exist.

Steps

1. Click **Management > Access Groups**.
2. Click the Actions icon for the access group you want to edit.
3. Click the **Edit** button.
4. To add a Fibre Channel initiator to the volume access group, perform the following steps:
 - a. Under Add Initiators, select an existing Fibre Channel initiator from the **Unbound Fibre Channel Initiators** list.
 - b. Click **Add FC Initiator**.

Note: You can create an initiator during this step if you click the **Create Initiator** link, enter an initiator name, and click **Create**. The system automatically adds the initiator to the **Initiators** list after you create it.

A sample of the format is as follows:

5f:47:ac:c0:5c:74:d4:02

5. To add an iSCSI initiator to the volume access group, under Add Initiators, select an existing initiator from the **Initiators** list.

Note: You can create an initiator during this step if you click the **Create Initiator** link, enter an initiator name, and click **Create**. The system automatically adds the initiator to the **Initiators** list after you create it.

The accepted format of an initiator IQN is as follows: iqn.yyyy-mm, in which y and m are digits, followed by text which must only contain digits, lower-case alphabetic characters, a period (.), colon (:) or dash (-).

A sample of the format is as follows:

iqn.2010-01.com.solidfire:c2r9.fc0.2100000e1e09bb8b

Tip: You can find the initiator IQN for each volume from the **Management > Volumes** Active Volumes page by clicking the Actions icon and then selecting **View Details** for the volume.

6. Click **Save Changes**.

Adding multiple initiators to a volume access group

You can add multiple initiators to an existing volume access group to allow access to volumes in the volume access group with or without requiring CHAP authentication..

About this task

When you add initiators to a volume access group, the initiators have access to all volumes in that volume access group.

Tip: You can find the initiator for each volume by clicking the Actions icon and then **View Details** for the volume in the active volumes list.

You can add multiple initiators to an existing volume access group to enable access to volumes and assign unique CHAP credentials for each initiator within that volume access group. This enables you to apply this option for volume access groups that already exist.

You can assign initiator-based CHAP attributes by using an API call. To add a CHAP account name and credentials per initiator, you must use the `ModifyInitiator` API call to remove and add CHAP access and attributes. For details, see the API reference information.

[Managing storage with the Element API](#)

Steps

1. Click **Management > Initiators**.
2. Select the initiators you want to add to an access group.
3. Click the **Bulk Actions** button.
4. Click **Add to Volume Access Group**.
5. In the Add to Volume Access Group dialog box, select an access group from the **Volume Access Group** list.
6. Click **Add**.

Removing initiators from an access group

When you remove an initiator from an access group, it can no longer access the volumes in that volume access group. Normal account access to the volume is not disrupted.

About this task

Modifying CHAP settings in an account or removing initiators or volumes from an access group can cause initiators to lose access to volumes unexpectedly. To verify that volume access will not be lost unexpectedly, always logout iSCSI sessions that will be affected by an account or access group change, and verify that initiators can reconnect to volumes after any changes to initiator settings and cluster settings have been completed.

Steps

1. Click **Management > Access Groups**.
2. Click the Actions icon for the access group you want to remove.
3. In the resulting menu, select **Edit**.
4. Under Add Initiators in the **Edit Volume Access Group** dialog box, click the arrow on the **Initiators** list.
5. Select the x icon for each initiator you want to remove from the access group.
6. Click **Save Changes**.

Deleting an access group

You can delete an access group when it is no longer needed. You do not need to delete Initiator IDs and Volume IDs from the volume access group before deleting the group. After you delete the access group, group access to the volumes is discontinued.

Steps

1. Click **Management > Access Groups**.
2. Click the Actions icon for the access group you want to delete.
3. In the resulting menu, click **Delete**.
4. To also delete the initiators associated with this access group, select the **Delete initiators in this access group** check box.
5. Confirm the action.

Deleting an initiator

You can delete an initiator after it is no longer needed. When you delete an initiator, the system removes it from any associated volume access group. Any connections using the initiator remain valid until the connection is reset.

Steps

1. Click **Management > Initiators**.
2. Perform the steps to delete a single initiator or multiple initiators:

Option	Steps
Delete single initiator	<ol style="list-style-type: none">a. Click the Actions icon for the initiator you want to delete.b. Click Delete.c. Confirm the action.
Delete multiple initiators	<ol style="list-style-type: none">a. Select the check boxes next to the initiators you want to delete.b. Click the Bulk Actions button.c. In the resulting menu, select Delete.d. Confirm the action.

Data protection

NetApp Element software enables you to protect your data in a variety of ways with capabilities such as snapshots for individual volumes or groups of volumes, replication between clusters and volumes running on Element, and replication to ONTAP systems.

Snapshots

Snapshot-only data protection replicates changed data at specific points of time to a remote cluster. Only those snapshots that are created on the source cluster are replicated. Active writes from the source volume are not.

Replication between clusters and volumes running on Element

You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair both running on Element for failover and failback scenarios.

Replication between Element and ONTAP clusters using SnapMirror technology

With NetApp SnapMirror technology, you can replicate snapshots that were taken using Element to ONTAP for disaster recovery purposes. In a SnapMirror relationship, Element is one endpoint and ONTAP is the other.

Related tasks

[Using volume snapshots for data protection](#) on page 97

A volume snapshot is a point-in-time copy of a volume. You can take a snapshot of a volume and use the snapshot later if you need to roll a volume back to the state it was in at the time the snapshot was created.

[Performing remote replication between clusters running NetApp Element software](#) on page 110

For clusters running Element software, real-time replication enables the quick creation of remote copies of volume data. You can pair a storage cluster with up to four other storage clusters. You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair for failover and failback scenarios.

[Using SnapMirror replication between Element and ONTAP clusters](#) on page 124

You can create SnapMirror relationships from the Data Protection tab in the NetApp Element UI. SnapMirror functionality must be enabled to see this in the user interface.

[Backing up and restoring volumes](#) on page 135

You can back up and restore volumes to other SolidFire storage, as well as secondary object stores that are compatible with Amazon S3 or OpenStack Swift.

Using volume snapshots for data protection

A volume snapshot is a point-in-time copy of a volume. You can take a snapshot of a volume and use the snapshot later if you need to roll a volume back to the state it was in at the time the snapshot was created.

About this task

Snapshots are similar to volume clones. However, snapshots are simply replicas of volume metadata, so you cannot mount or write to them. Creating a volume snapshot also takes only a small amount of system resources and space, which makes snapshot creation faster than cloning.

You can take a snapshot of an individual volume or a set of volumes.

Optionally, replicate snapshots to a remote cluster and use them as a backup copy of the volume. This enables you to roll back a volume to a specific point in time by using the replicated snapshot. Alternatively, you can create a clone of a volume from a replicated snapshot.

Related tasks

[Using individual volume snapshots for data protection task](#) on page 98

A volume snapshot is a point-in-time copy of a volume. You can use an individual volume rather than a group of volumes for the snapshot.

[Using group snapshots for data protection task](#) on page 102

You can create a group snapshot of a related set of volumes to preserve a point-in-time copy of the metadata for each volume. You can use the group snapshot in the future as a backup or rollback to restore the state of the group of volumes to a previous state.

[Scheduling a snapshot](#) on page 107

You can protect data on a volume or a group of volumes by scheduling volume snapshots to occur at specified intervals. You can schedule either single volume snapshots or group snapshots to run automatically.

Using individual volume snapshots for data protection task

A volume snapshot is a point-in-time copy of a volume. You can use an individual volume rather than a group of volumes for the snapshot.

Related tasks

[Creating a volume snapshot](#) on page 99

You can create a snapshot of an active volume to preserve the volume image at any point in time. You can create up to 32 snapshots for a single volume.

[Editing snapshot retention](#) on page 99

You can change the retention period for a snapshot to control when or if the system deletes snapshots. The retention period you specify begins when you enter the new interval. When you set a retention period, you can select a period that begins at the current time (retention is not calculated from the snapshot creation time). You can specify intervals in minutes, hours, and days.

[Deleting a snapshot](#) on page 100

You can delete a volume snapshot from a storage cluster running Element software. When you delete a snapshot, the system immediately removes it.

[Cloning a volume from a snapshot](#) on page 100

You can create a new volume from a snapshot of a volume. When you do this, the system uses the snapshot information to clone a new volume using the data contained on the volume at the time the snapshot was created. This process stores information about other snapshots of the volume in the newly created volume.

[Rolling back a volume to a snapshot](#) on page 100

You can roll back a volume to a previous snapshot at any time. This reverts any changes made to the volume since the snapshot was created.

[Backing up a volume snapshot to an Amazon S3 object store](#) on page 101

You can back up SolidFire snapshots to external object stores that are compatible with Amazon S3.

[Backing up a volume snapshot to an OpenStack Swift object store](#) on page 101

You can back up SolidFire snapshots to secondary object stores that are compatible with OpenStack Swift.

[Backing up a volume snapshot to a SolidFire cluster](#) on page 102

You can back up volume snapshots residing on a SolidFire cluster to a remote SolidFire cluster.

Creating a volume snapshot

You can create a snapshot of an active volume to preserve the volume image at any point in time. You can create up to 32 snapshots for a single volume.

Steps

1. Click **Management > Volumes**.
2. Click the Actions icon for the volume you want to use for the snapshot.
3. In the resulting menu, select **Snapshot**.
4. In the **Create Snapshot of Volume** dialog box, enter the new snapshot name.
5. Optional: Select the **Include Snapshot in Replication When Paired** check box to ensure that the snapshot is captured in replication when the parent volume is paired.
6. To set the retention for the snapshot, select from one of the following options:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.
7. To take a single, immediate snapshot, perform the following steps:
 - a. Click **Take Snapshot Now**.
 - b. Click **Create Snapshot**.
8. To schedule the snapshot to run at a future time, perform the following steps:
 - a. Click **Create Snapshot Schedule**.
 - b. Enter a **New Schedule Name**.
 - c. Choose a **Schedule Type** from the list.
 - d. Optional: Select the **Recurring Schedule** check box to repeat the scheduled snapshot periodically.
 - e. Click **Create Schedule**.

Related tasks

[Scheduling a snapshot](#) on page 107

You can protect data on a volume or a group of volumes by scheduling volume snapshots to occur at specified intervals. You can schedule either single volume snapshots or group snapshots to run automatically.

Editing snapshot retention

You can change the retention period for a snapshot to control when or if the system deletes snapshots. The retention period you specify begins when you enter the new interval. When you set a retention period, you can select a period that begins at the current time (retention is not calculated from the snapshot creation time). You can specify intervals in minutes, hours, and days.

Steps

1. Click **Data Protection > Snapshots**.
2. Click the Actions icon for the snapshot you want to edit.
3. In the resulting menu, click **Edit**.
4. Optional: Select the **Include Snapshot in Replication When Paired** check box to ensure that the snapshot is captured in replication when the parent volume is paired.
5. Optional: Select a retention option for the snapshot:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.

- Click **Set Retention Period** and use the date spin boxes to select a length of time for the system to retain the snapshot.

6. Click **Save Changes**.

Deleting a snapshot

You can delete a volume snapshot from a storage cluster running Element software. When you delete a snapshot, the system immediately removes it.

About this task

You can delete snapshots that are being replicated from the source cluster. If a snapshot is syncing to the target cluster when you delete it, the sync replication completes and the snapshot is deleted from the source cluster. The snapshot is not deleted from the target cluster.

You can also delete snapshots that have been replicated to the target from the target cluster. The deleted snapshot is kept in a list of deleted snapshots on the target until the system detects that you have deleted the snapshot on the source cluster. When the target detects that you have deleted the source snapshot, the target stops replication of the snapshot.

When you delete a snapshot from the source cluster, the target cluster snapshot is not affected (the reverse is also true).

Steps

1. Click **Data Protection > Snapshots**.
2. Click the Actions icon for the snapshot you want to delete.
3. In the resulting menu, select **Delete**.
4. Confirm the action.

Cloning a volume from a snapshot

You can create a new volume from a snapshot of a volume. When you do this, the system uses the snapshot information to clone a new volume using the data contained on the volume at the time the snapshot was created. This process stores information about other snapshots of the volume in the newly created volume.

Steps

1. Click **Data Protection > Snapshots**.
2. Click the Actions icon for the snapshot you want to use for the volume clone.
3. In the resulting menu, click **Clone Volume From Snapshot**.
4. Enter a **Volume Name** in the **Clone Volume From Snapshot** dialog box.
5. Select a **Total Size** and size units for the new volume.
6. Select an **Access** type for the volume.
7. Select an **Account** from the list to associate with the new volume.
8. Click **Start Cloning**.

Rolling back a volume to a snapshot

You can roll back a volume to a previous snapshot at any time. This reverts any changes made to the volume since the snapshot was created.

Steps

1. Click **Data Protection > Snapshots**.
2. Click the Actions icon for the snapshot you want to use for the volume rollback.
3. In the resulting menu, select **Rollback Volume To Snapshot**.

4. Optional: To save the current state of the volume before rolling back to the snapshot:
 - a. In the **Rollback To Snapshot** dialog box, select **Save volume's current state as a snapshot**.
 - b. Enter a name for the new snapshot.
5. Click **Rollback Snapshot**.

Volume snapshot backup operations

You can use the integrated backup feature to back up a volume snapshot. You can back up snapshots from a SolidFire cluster to an external object store, or to another SolidFire cluster. When you back up a snapshot to an external object store, you must have a connection to the object store that allows read/write operations.

Backing up a volume snapshot to an Amazon S3 object store

You can back up SolidFire snapshots to external object stores that are compatible with Amazon S3.

Steps

1. Click **Data Protection > Snapshots**.
2. Click the Actions icon for the snapshot you want to back up.
3. In the resulting menu, click **Backup to**.
4. In the **Integrated Backup** dialog box under **Backup to**, select **S3**.
5. Select an option under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Enter a hostname to use to access the object store in the **Hostname** field.
7. Enter an access key ID for the account in the **Access Key ID** field.
8. Enter the secret access key for the account in the **Secret Access Key** field.
9. Enter the S3 bucket in which to store the backup in the **S3 Bucket** field.
10. Optional: Enter a nametag to append to the prefix in the **Nametag** field.
11. Click **Start Read**.

Backing up a volume snapshot to an OpenStack Swift object store

You can back up SolidFire snapshots to secondary object stores that are compatible with OpenStack Swift.

Steps

1. Click **Data Protection > Snapshots**.
2. Click the Actions icon for the snapshot you want to back up.
3. In the resulting menu, click **Backup to**.
4. In the **Integrated Backup** dialog box, under **Backup to**, select **Swift**.
5. Select an option under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Enter a **URL** to use to access the object store.
7. Enter a **Username** for the account.
8. Enter the **Authentication Key** for the account.
9. Enter the **Container** in which to store the backup.
10. Optional: Enter a **Nametag**.

11. Click Start Read.

Backing up a volume snapshot to a SolidFire cluster

You can back up volume snapshots residing on a SolidFire cluster to a remote SolidFire cluster.

Before you begin

Ensure that the source and target clusters are paired.

About this task

When backing up or restoring from one cluster to another, the system generates a key to be used as authentication between the clusters. This bulk volume write key allows the source cluster to authenticate with the destination cluster, providing a level of security when writing to the destination volume. As part of the backup or restore process, you need to generate a bulk volume write key from the destination volume before starting the operation.

Steps

1. On the destination cluster, click **Management > Volumes**.
2. Click the Actions icon for the destination volume.
3. In the resulting menu, click **Restore from**.
4. In the **Integrated Restore** dialog box under **Restore from**, select **SolidFire**.
5. Select a data format under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Click **Generate Key**.
7. Copy the key from the **Bulk Volume Write Key** box to your clipboard.
8. On the source cluster, click **Data Protection > Snapshots**.
9. Click the Actions icon for the snapshot you want to use for the backup.
10. In the resulting menu, click **Backup to**.
11. In the **Integrated Backup** dialog box under **Backup to**, select **SolidFire**.
12. Select the same data format you selected earlier in the **Data Format** field.
13. Enter the management virtual IP address of the destination volume's cluster in the **Remote Cluster MVIP** field.
14. Enter the remote cluster user name in the **Remote Cluster Username** field.
15. Enter the remote cluster password in the **Remote Cluster Password** field.
16. In the **Bulk Volume Write Key** field, paste the key you generated on the destination cluster earlier.
17. Click **Start Read**.

Using group snapshots for data protection task

You can create a group snapshot of a related set of volumes to preserve a point-in-time copy of the metadata for each volume. You can use the group snapshot in the future as a backup or rollback to restore the state of the group of volumes to a previous state.

Related tasks

[Creating a group snapshot](#) on page 104

You can create a snapshot of a group of volumes, and you can also create a group snapshot schedule to automate group snapshots. A single group snapshot can consistently snapshot up to 32 volumes at one time.

[Editing group snapshots](#) on page 104

You can edit the replication and retention settings for existing group snapshots.

[Editing members of group snapshot](#) on page 105

You can edit the retention settings for members of an existing group snapshot.

[Deleting a group snapshot](#) on page 105

You can delete a group snapshot from the system. When you delete the group snapshot, you can choose whether all snapshots associated with the group are deleted or retained as individual snapshots.

[Rolling back volumes to a group snapshot](#) on page 105

You can roll back a group of volumes at any time to a group snapshot.

[Cloning multiple volumes](#) on page 106

You can create multiple volume clones in a single operation to create a point-in-time copy of the data on a group of volumes.

[Cloning multiple volumes from a group snapshot](#) on page 106

You can clone a group of volumes from a point-in-time group snapshot. This operation requires that a group snapshot of the volumes already exist, because the group snapshot is used as the basis to create the volumes. After you create the volumes, you can use them like any other volume in the system.

Group snapshot details

The Group Snapshots page on the Data Protection tab provides information about the group snapshots.

ID

The system-generated ID for the group snapshot.

UUID

The unique ID of the group snapshot.

Name

User-defined name for the group snapshot.

Create Time

The time at which the group snapshot was created.

Status

The current status of the snapshot. Possible values:

- **Preparing:** The snapshot is being prepared for use and is not yet writable.
- **Done:** This snapshot has finished preparation and is now usable.
- **Active:** The snapshot is the active branch.

Volumes

The number of volumes in the group.

Retain Until

The day and time the snapshot will be deleted.

Remote Replication

Indication of whether or not the snapshot is enabled for replication to a remote SolidFire cluster. Possible values:

- **Enabled:** The snapshot is enabled for remote replication.
- **Disabled:** The snapshot is not enabled for remote replication.

Creating a group snapshot

You can create a snapshot of a group of volumes, and you can also create a group snapshot schedule to automate group snapshots. A single group snapshot can consistently snapshot up to 32 volumes at one time.

Steps

1. Click **Management > Volumes**.
2. Use the check boxes to select multiple volumes for a group of volumes.
3. Click **Bulk Actions**.
4. Click **Group Snapshot**.
5. Enter a new group snapshot name in the Create Group Snapshot of Volumes dialog box.
6. Optional: Select the **Include Each Group Snapshot Member in Replication When Paired** check box to ensure that each snapshot is captured in replication when the parent volume is paired.
7. Select a retention option for the group snapshot:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.
8. To take a single, immediate snapshot, perform the following steps:
 - a. Click **Take Group Snapshot Now**.
 - b. Click **Create Group Snapshot**.
9. To schedule the snapshot to run at a future time, perform the following steps:
 - a. Click **Create Group Snapshot Schedule**.
 - b. Enter a **New Schedule Name**.
 - c. Select a **Schedule Type** from the list.
 - d. Optional: Select the **Recurring Schedule** check box to repeat the scheduled snapshot periodically.
 - e. Click **Create Schedule**.

Editing group snapshots

You can edit the replication and retention settings for existing group snapshots.

Steps

1. Click **Data Protection > Group Snapshots**.
2. Click the Actions icon for the group snapshot you want to edit.
3. In the resulting menu, select **Edit**.
4. Optional: To change the replication setting for the group snapshot:
 - a. Click **Edit** next to **Current Replication**.
 - b. Select the **Include Each Group Snapshot Member in Replication When Paired** check box to ensure that each snapshot is captured in replication when the parent volume is paired.
5. Optional: To change the retention setting for the group snapshot, select from the following options:
 - a. Click **Edit** next to **Current Retention**.
 - b. Select a retention option for the group snapshot:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.

6. Click **Save Changes.**

Deleting a group snapshot

You can delete a group snapshot from the system. When you delete the group snapshot, you can choose whether all snapshots associated with the group are deleted or retained as individual snapshots.

About this task

If you delete a volume or snapshot that is a member of a group snapshot, you can no longer roll back to the group snapshot. However, you can roll back each volume individually.

Steps

1. Click **Data Protection > Group Snapshots**.
2. Click the Actions icon for the snapshot you want to delete.
3. In the resulting menu, click **Delete**.
4. Select from one of the following options in the confirmation dialog box:
 - Click **Delete group snapshot AND all group snapshot members** to delete the group snapshot and all member snapshots.
 - Click **Retain group snapshot members as individual snapshots** to delete the group snapshot but keep all member snapshots.
5. Confirm the action.

Rolling back volumes to a group snapshot

You can roll back a group of volumes at any time to a group snapshot.

About this task

When you roll back a group of volumes, all volumes in the group are restored to the state they were in at the time the group snapshot was created. Rolling back also restores volume sizes to the size recorded in the original snapshot. If the system has purged a volume, all snapshots of that volume were also deleted at the time of the purge; the system does not restore any deleted volume snapshots.

Steps

1. Click **Data Protection > Group Snapshots**.
2. Click the Actions icon for the group snapshot you want to use for the volume rollback.
3. In the resulting menu, select **Rollback Volumes To Group Snapshot**.
4. Optional: To save the current state of the volumes before rolling back to the snapshot:
 - a. In the **Rollback To Snapshot** dialog box, select **Save volumes' current state as a group snapshot**.
 - b. Enter a name for the new snapshot.
5. Click **Rollback Group Snapshot**.

Editing members of group snapshot

You can edit the retention settings for members of an existing group snapshot.

Steps

1. Click **Data Protection > Snapshots**.
2. Click the **Members** tab.
3. Click the Actions icon for the group snapshot member you want to edit.
4. In the resulting menu, select **Edit**.

5. To change the replication setting for the snapshot, select from the following options:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.
6. Click **Save Changes**.

Cloning multiple volumes

You can create multiple volume clones in a single operation to create a point-in-time copy of the data on a group of volumes.

About this task

When you clone a volume, the system creates a snapshot of the volume and then creates a new volume from the data in the snapshot. You can mount and write to the new volume clone. Cloning multiple volumes is an asynchronous process and takes a variable amount of time depending on the size and number of the volumes being cloned.

Volume size and current cluster load affect the time needed to complete a cloning operation.

Steps

1. Click **Management > Volumes**.
2. Click the **Active** tab.
3. Use the check boxes to select multiple volumes, creating a group of volumes.
4. Click **Bulk Actions**.
5. Click **Clone** in the resulting menu.
6. Enter a **New Volume Name Prefix** in the **Clone Multiple Volumes** dialog box.
The prefix is applied to all volumes in the group.
7. Optional: Select a different account to which the clone will belong.
If you do not select an account, the system assigns the new volumes to the current volume account.
8. Optional: Select a different access method for the volumes in the clone.
If you do not select an access method, the system uses the current volume access.
9. Click **Start Cloning**.

Cloning multiple volumes from a group snapshot

You can clone a group of volumes from a point-in-time group snapshot. This operation requires that a group snapshot of the volumes already exist, because the group snapshot is used as the basis to create the volumes. After you create the volumes, you can use them like any other volume in the system.

About this task

Volume size and current cluster load affect the time needed to complete a cloning operation.

Steps

1. Click **Data Protection > Group Snapshots**.
2. Click the Actions icon for the group snapshot you want to use for the volume clones.
3. In the resulting menu, select **Clone Volumes From Group Snapshot**.
4. Enter a **New Volume Name Prefix** in the **Clone Volumes From Group Snapshot** dialog box.
The prefix is applied to all volumes created from the group snapshot.
5. Optional: Select a different account to which the clone will belong.

If you do not select an account, the system assigns the new volumes to the current volume account.

6. Optional: Select a different access method for the volumes in the clone.

If you do not select an access method, the system uses the current volume access.

7. Click **Start Cloning**.

Scheduling a snapshot

You can protect data on a volume or a group of volumes by scheduling volume snapshots to occur at specified intervals. You can schedule either single volume snapshots or group snapshots to run automatically.

About this task

When you configure a snapshot schedule, you can choose from time intervals based on days of the week or days of the month. You can also specify the days, hours, and minutes before the next snapshot occurs. You can store the resulting snapshots on a remote storage system if the volume is being replicated.

Related tasks

[Creating a snapshot schedule](#) on page 108

You can schedule a snapshot of a volume or volumes to automatically occur at specified intervals.

[Editing a snapshot schedule](#) on page 108

You can modify existing snapshot schedules. After modification, the next time the schedule runs it uses the updated attributes. Any snapshots created by the original schedule remain on the storage system.

[Deleting a snapshot schedule](#) on page 109

You can delete a snapshot schedule. After you delete the schedule, it does not run any future scheduled snapshots. Any snapshots that were created by the schedule remain on the storage system.

[Copying a snapshot schedule](#) on page 109

You can copy a schedule and maintain its current attributes.

Snapshot schedule details

On the Data Protection > Schedules page, you can view the following information in the list of snapshot schedules.

ID

The system-generated ID for the snapshot.

Type

The type of schedule. Snapshot is currently the only type supported.

Name

The name given to the schedule when it was created. Snapshot schedule names can be up to 223 characters in length and contain a-z, 0-9, and dash (-) characters.

Frequency

The frequency at which the schedule is run. The frequency can be set in hours and minutes, weeks, or months.

Recurring

Indication of whether the schedule is to run only once or at regular intervals.

Manually Paused

Indication of whether or not the schedule has been manually paused.

Volume IDs

The ID of the volume the schedule will use when the schedule is run.

Last Run

The last time the schedule was run.

Last Run Status

The outcome of the last schedule execution. Possible values:

- Success
- Failure

Creating a snapshot schedule

You can schedule a snapshot of a volume or volumes to automatically occur at specified intervals.

About this task

When you configure a snapshot schedule, you can choose from time intervals based on days of the week or days of the month. You can also create a recurring schedule and specify the days, hours, and minutes before the next snapshot occurs.

If you schedule a snapshot to run at a time period that is not divisible by 5 minutes, the snapshot will run at the next time period that is divisible by 5 minutes. For example, if you schedule a snapshot to run at 12:42:00 UTC, it will run at 12:45:00 UTC. You cannot schedule a snapshot to run at intervals of less than 5 minutes.

Steps

1. Click **Data Protection > Schedules**.
2. Click **Create Schedule**.
3. In the **Volume IDs CSV** field, enter a single volume ID or a comma-separated list of volume IDs to include in the snapshot operation.
4. Enter a new schedule name.
5. Select a schedule type and set the schedule from the options provided.
6. Optional: Select **Recurring Schedule** to repeat the snapshot schedule indefinitely.
7. Optional: Enter a name for the new snapshot in the **New Snapshot Name** field. If you leave the field blank, the system uses the time and date of the snapshot's creation as the name.
8. Optional: Select the **Include Snapshots in Replication When Paired** check box to ensure that the snapshots are captured in replication when the parent volume is paired.
9. To set the retention for the snapshot, select from the following options:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to choose a length of time for the system to retain the snapshot.
10. Click **Create Schedule**.

Editing a snapshot schedule

You can modify existing snapshot schedules. After modification, the next time the schedule runs it uses the updated attributes. Any snapshots created by the original schedule remain on the storage system.

Steps

1. Click **Data Protection > Schedules**.
2. Click the Actions icon for the schedule you want to change.

3. In the resulting menu, click **Edit**.
4. In the **Volume IDs CSV** field, modify the single volume ID or comma-separated list of volume IDs currently included in the snapshot operation.
5. To pause or resume the schedule, select from the following options:
 - To pause an active schedule, select **Yes** from the **Manually Pause Schedule** list.
 - To resume a paused schedule, select **No** from the **Manually Pause Schedule** list.
6. Enter a different name for the schedule in the **New Schedule Name** field if desired.
7. To change the schedule to run on different days of the week or month, select **Schedule Type** and change the schedule from the options provided.
8. Optional: Select **Recurring Schedule** to repeat the snapshot schedule indefinitely.
9. Optional: Enter or modify the name for the new snapshot in the **New Snapshot Name** field.
If you leave the field blank, the system uses the time and date of the snapshot's creation as the name.
10. Optional: Select the **Include Snapshots in Replication When Paired** check box to ensure that the snapshots are captured in replication when the parent volume is paired.
11. To change the retention setting, select from the following options:
 - Click **Keep Forever** to retain the snapshot on the system indefinitely.
 - Click **Set Retention Period** and use the date spin boxes to select a length of time for the system to retain the snapshot.
12. Click **Save Changes**.

Copying a snapshot schedule

You can copy a schedule and maintain its current attributes.

Steps

1. Click **Data Protection > Schedules**.
2. Click the Actions icon for the schedule you want to copy.
3. In the resulting menu, click **Make a Copy**.

The **Create Schedule** dialog box appears, populated with the current attributes of the schedule.

4. Optional: Enter a name and updated attributes for the new schedule.
5. Click **Create Schedule**.

Deleting a snapshot schedule

You can delete a snapshot schedule. After you delete the schedule, it does not run any future scheduled snapshots. Any snapshots that were created by the schedule remain on the storage system.

Steps

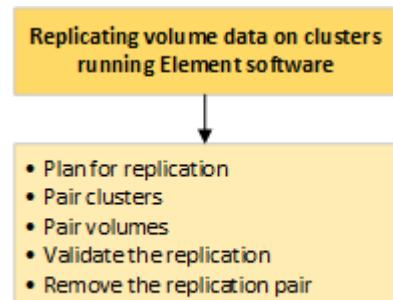
1. Click **Data Protection > Schedules**.
2. Click the Actions icon for the schedule you want to delete.
3. In the resulting menu, click **Delete**.
4. Confirm the action.

Performing remote replication between clusters running NetApp Element software

For clusters running Element software, real-time replication enables the quick creation of remote copies of volume data. You can pair a storage cluster with up to four other storage clusters. You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair for failover and failback scenarios.

About this task

The replication process includes these steps:



Steps

1. [Planning cluster and volume pairing for real-time replication](#) on page 111

Real-time remote replication requires that you pair two storage clusters running Element software, pair volumes on each cluster, and validate replication. After replication completes, you should delete the volume relationship.

2. [Pairing clusters](#) on page 111

You must pair two clusters as a first step to using real-time replication functionality. After you pair and connect two clusters, you can configure active volumes on one cluster to be continuously replicated to a second cluster, providing continuous data protection (CDP).

3. [Pairing volumes](#) on page 114

After you have established a connection between clusters in a cluster pair, you can pair a volume on one cluster with a volume on the other cluster in the pair. When a volume pairing relationship is established, you must identify which volume is the replication target.

4. [Validating volume replication](#) on page 119

After a volume is replicated, you should ensure that the source and target volumes are active. When in an active state, volumes are paired, data is being sent from the source to the target volume, and the data is in sync.

5. [Deleting a volume relationship after replication](#) on page 119

After replication completes and you no longer need the volume pair relationship, you can delete the volume relationship.

6. [Managing volume relationships](#) on page 120

You can manage volume relationships in many ways, such as pausing replication, reversing volume pairing, changing the mode of replication, deleting a volume pair, or deleting a cluster pair.

Planning cluster and volume pairing for real-time replication

Real-time remote replication requires that you pair two storage clusters running Element software, pair volumes on each cluster, and validate replication. After replication completes, you should delete the volume relationship.

Before you begin

- You must have cluster administrator privileges to one or both clusters being paired.
- All node IP addresses on both management and storage networks for paired clusters are routed to each other.
- MTU of all paired nodes must be the same and be supported end-to-end between clusters.
- Both storage clusters should have unique cluster names, MVIPs, SVIPs., and all node IP addresses.
- The difference between Element software versions on the clusters is no greater than one major version. If the difference is greater, one of the clusters must be upgraded to perform data replication.

Note: WAN Accelerator appliances have not been qualified by NetApp for use when replicating data. These appliances can interfere with compression and deduplication if deployed between two clusters that are replicating data. Be sure to fully qualify the effects of any WAN Accelerator appliance before you deploy it in a production environment.

Related tasks

[Pairing clusters](#) on page 111

You must pair two clusters as a first step to using real-time replication functionality. After you pair and connect two clusters, you can configure active volumes on one cluster to be continuously replicated to a second cluster, providing continuous data protection (CDP).

[Pairing volumes](#) on page 114

After you have established a connection between clusters in a cluster pair, you can pair a volume on one cluster with a volume on the other cluster in the pair. When a volume pairing relationship is established, you must identify which volume is the replication target.

[Assigning a replication source and target to paired volumes](#) on page 118

After volumes are paired, you must assign a source volume and its replication target volume. A replication source or target can be either volume in a volume pair. You can also use this procedure to redirect data sent to a source volume to a remote target volume should the source volume become unavailable.

Pairing clusters

You must pair two clusters as a first step to using real-time replication functionality. After you pair and connect two clusters, you can configure active volumes on one cluster to be continuously replicated to a second cluster, providing continuous data protection (CDP).

Before you begin

- You must have cluster administrator privileges to one or both clusters being paired.
- All node MIPs and SIPs are routed to each other.
- Less than 2000 ms of round-trip latency between clusters.
- Both storage clusters should have unique cluster names, MVIPs, SVIPs, and all node IP addresses.

- The difference between Element software versions on the clusters is no greater than one major version. If the difference is greater, one of the clusters must be upgraded to perform data replication.

Note: Cluster pairing requires full connectivity between nodes on the management network. Replication requires connectivity between the individual nodes on the storage cluster network.

About this task

You can pair one cluster with up to four other clusters for replicating volumes. You can also pair clusters within the cluster group with each other.

Related reference

[Network port requirements](#) on page 13

You might need to allow the following TCP ports through your datacenter's edge firewall so that you can manage the system remotely and allow clients outside of your datacenter to connect to resources. Some of these ports might not be required, depending on how you use the system.

Steps

1. [Pairing clusters using MVIP or a pairing key](#) on page 112

You can pair a source and target cluster using the MVIP of the target cluster if there is cluster administrator access to both clusters. If cluster administrator access is only available on one cluster in a cluster pair, a pairing key can be used on the target cluster to complete the cluster pairing.

2. [Validating the cluster pair connection](#) on page 114

After the cluster pairing has completed, you might want to verify the cluster pair connection to ensure replication success.

Pairing clusters using MVIP or a pairing key

You can pair a source and target cluster using the MVIP of the target cluster if there is cluster administrator access to both clusters. If cluster administrator access is only available on one cluster in a cluster pair, a pairing key can be used on the target cluster to complete the cluster pairing.

Step

Select one of the following methods to pair clusters:

- Pair clusters using MVIP: Use this method if there is cluster administrator access to both clusters. This method uses the MVIP of the remote cluster to pair two clusters.
- Pair clusters using a pairing key: Use this method if there is cluster administrator access to only one of the clusters. This method generates a pairing key that can be used on the target cluster to complete the cluster pairing.

Related tasks

[Pairing clusters using MVIP](#) on page 113

You can pair two clusters for real-time replication by using the MVIP of one cluster to establish a connection with the other cluster. Cluster administrator access on both of clusters is required to use this method. The cluster administrator user name and password is used to authenticate cluster access before the clusters can be paired.

[Pairing clusters using a pairing key](#) on page 113

If you have cluster administrator access to a local cluster but not the remote cluster, you can pair the clusters using a pairing key. A pairing key is generated on a local cluster and then sent securely to a cluster administrator at a remote site to establish a connection and complete the cluster pairing for real-time replication.

Pairing clusters using MVIP

You can pair two clusters for real-time replication by using the MVIP of one cluster to establish a connection with the other cluster. Cluster administrator access on both of clusters is required to use this method. The cluster administrator user name and password is used to authenticate cluster access before the clusters can be paired.

Steps

1. On the local cluster, select **Data Protection > Cluster Pairs**.
2. Click **Pair Cluster**.
3. Click **Start Pairing** and click **Yes** to indicate that you have access to the remote cluster.
4. Enter the remote cluster MVIP address.
5. Click **Complete pairing on remote cluster**.
In the **Authentication Required** window, enter the cluster administrator user name and password of the remote cluster.
6. On the remote cluster, select **Data Protection > Cluster Pairs**.
7. Click **Pair Cluster**.
8. Click **Complete Pairing**.
9. Click the **Complete Pairing** button.

Related tasks

[Pairing clusters using a pairing key](#) on page 113

If you have cluster administrator access to a local cluster but not the remote cluster, you can pair the clusters using a pairing key. A pairing key is generated on a local cluster and then sent securely to a cluster administrator at a remote site to establish a connection and complete the cluster pairing for real-time replication.

Related information

[Pairing clusters using MVIP \(video\)](#)

Pairing clusters using a pairing key

If you have cluster administrator access to a local cluster but not the remote cluster, you can pair the clusters using a pairing key. A pairing key is generated on a local cluster and then sent securely to a cluster administrator at a remote site to establish a connection and complete the cluster pairing for real-time replication.

Steps

1. On the local cluster, select **Data Protection > Cluster Pairs**.
2. Click **Pair Cluster**.
3. Click **Start Pairing** and click **No** to indicate that you do not have access to the remote cluster.
4. Click **Generate Key**.

Note: This action generates a text key for pairing and creates an unconfigured cluster pair on the local cluster. If you do not complete the procedure, you will need to manually delete the cluster pair.

5. Copy the cluster pairing key to your clipboard.
6. Make the pairing key accessible to the cluster administrator at the remote cluster site.

Note: The cluster pairing key contains a version of the MVIP, user name, password, and database information to permit volume connections for remote replication. This key should be treated in a secure manner and not stored in a way that would allow accidental or unsecured access to the user name or password.



Attention: Do not modify any of the characters in the pairing key. The key becomes invalid if it is modified.

7. On the remote cluster, select **Data Protection > Cluster Pairs**.
8. Click **Pair Cluster**.
9. Click **Complete Pairing** and enter the pairing key in the **Pairing Key** field (paste is the recommended method).
10. Click **Complete Pairing**.

Related tasks

[Pairing clusters using MVIP](#) on page 113

You can pair two clusters for real-time replication by using the MVIP of one cluster to establish a connection with the other cluster. Cluster administrator access on both of clusters is required to use this method. The cluster administrator user name and password is used to authenticate cluster access before the clusters can be paired.

Related information

[Pairing clusters using a cluster pairing key \(video\)](#)

Validating the cluster pair connection

After the cluster pairing has completed, you might want to verify the cluster pair connection to ensure replication success.

Steps

1. On the local cluster, select **Data Protection > Cluster Pairs**.
2. In the **Cluster Pairs** window, verify that the cluster pair is connected.
3. Optional: Navigate back to the local cluster and the **Cluster Pairs** window and verify that the cluster pair is connected.

Pairing volumes

After you have established a connection between clusters in a cluster pair, you can pair a volume on one cluster with a volume on the other cluster in the pair. When a volume pairing relationship is established, you must identify which volume is the replication target.

Before you begin

- You have established a connection between clusters in a cluster pair.
- You have cluster administrator privileges to one or both clusters being paired.

About this task

You can pair two volumes for real-time replication that are stored on different storage clusters in a connected cluster pair. After you pair two clusters, you can configure active volumes on one cluster to be continuously replicated to a second cluster, providing continuous data protection (CDP). You can also assign either volume to be the source or target of the replication.

Volume pairings are always one-to-one. After a volume is part of a pairing with a volume on another cluster, you cannot pair it again with any other volume.

Steps

1. [Creating a target volume with read/write access](#) on page 115

The replication process involves two endpoints: the source and the target volume. When you create the target volume, the volume is automatically set to read/write mode to accept the data during the replication.

2. *Pairing volumes using a volume ID or pairing key* on page 115

The pairing process involves pairing two volumes by using either a volume ID or a pairing key.

3. *Assigning a replication source and target to paired volumes* on page 118

After volumes are paired, you must assign a source volume and its replication target volume. A replication source or target can be either volume in a volume pair. You can also use this procedure to redirect data sent to a source volume to a remote target volume should the source volume become unavailable.

Creating a target volume with read/write access

The replication process involves two endpoints: the source and the target volume. When you create the target volume, the volume is automatically set to read/write mode to accept the data during the replication.

About this task

Steps

1. Select **Management > Volumes**.
2. Click **Create Volume**.
3. In the Create a New Volume dialog box, enter the **Volume Name**.
4. Enter the total size of the volume, select a block size for the volume, and select the account that should have access to the volume.
5. Click **Create Volume**.
6. In the Active window, click the Actions icon for the volume.
7. Click **Edit**.
8. Change the account access level to **Replication Target**.
9. Click **Save Changes**.

Pairing volumes using a volume ID or pairing key

The pairing process involves pairing two volumes by using either a volume ID or a pairing key.

Step

Pair volumes by selecting one of the following methods:

- Using a volume ID: Use this method if you have cluster administrator access to both clusters on which volumes are to be paired. This method uses the volume ID of the volume on the remote cluster to initiate a connection.
- Using a pairing Key: Use this method if you have cluster administrator access to only the source cluster. This method generates a pairing key that can be used on the remote cluster to complete the volume pair.

Note: The volume pairing key contains an encrypted version of the volume information and might contain sensitive information. Only share this key in a secure manner.

Related tasks

Pairing volumes using a volume ID on page 116

You can pair a volume with another volume on a remote cluster if you have cluster administrator credentials for the remote cluster.

Pairing volumes using a pairing key on page 117

If you do not have cluster admin credentials for a remote cluster, you can pair a volume with another volume on a remote cluster using a pairing key.

Pairing volumes using a volume ID

You can pair a volume with another volume on a remote cluster if you have cluster administrator credentials for the remote cluster.

Before you begin

- Ensure that the clusters containing the volumes are paired.
- Create a new volume on the remote cluster.

Note: You can assign a replication source and target after the pairing process. A replication source or target can be either volume in a volume pair. You should create a target volume that contains no data and has the exact characteristics of the source volume, such as size, block size setting for the volumes (either 512e or 4k), and QoS configuration. If you assign an existing volume as the replication target, the data on that volume will be overwritten. The target volume can be greater or equal in size to the source volume, but it cannot be smaller.

- Know the target Volume ID.

Steps

1. Select **Management > Volumes**.
2. Click the Actions icon for the volume you want to pair.
3. Click **Pair**.
4. In the **Pair Volume** dialog box, select **Start Pairing**.
5. Select **I Do** to indicate that you have access to the remote cluster.
6. Select a **Replication Mode** from the list:
 - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
 - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
 - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.
7. Select a remote cluster from the list.
8. Choose a remote volume ID.
9. Click **Start Pairing**.

The system opens a web browser tab that connects to the Element UI of the remote cluster. You might be required to log on to the remote cluster with cluster administrator credentials.

10. In the Element UI of the remote cluster, select **Complete Pairing**.
11. Confirm the details in **Confirm Volume Pairing**.
12. Click **Complete Pairing**.

After you confirm the pairing, the two clusters begin the process of connecting the volumes for pairing. During the pairing process, you can see messages in the **Volume Status** column of the **Volume Pairs** window. The volume pair displays the following message until the volume pair source and target are assigned: `PausedMisconfigured`

Related tasks

[Assigning a replication source and target to paired volumes](#) on page 118

After volumes are paired, you must assign a source volume and its replication target volume. A replication source or target can be either volume in a volume pair. You can also use this procedure

to redirect data sent to a source volume to a remote target volume should the source volume become unavailable.

Related reference

[Volume pairing messages](#) on page 122

You can view volume pairing messages during the initial pairing process from the Volume Pairs page under the Data Protection tab. These messages can display on both source and target ends of the pair in the Replicating Volumes list view.

[Volume pairing warnings](#) on page 123

The Volume Pairs page on the Data Protection tab provides these messages after you pair volumes. These messages can display on both source and target ends of the pair (unless otherwise indicated) in the Replicating Volumes list view.

Pairing volumes using a pairing key

If you do not have cluster admin credentials for a remote cluster, you can pair a volume with another volume on a remote cluster using a pairing key.

Before you begin

- Ensure that the clusters containing the volumes are paired.
- Ensure that there is a volume on the remote cluster to use for the pairing.

Note: You can assign a replication source and target after the pairing process. A replication source or target can be either volume in a volume pair. You should create a target volume that contains no data and has the exact characteristics of the source volume, such as size, block size setting for the volumes (either 512e or 4k), and QoS configuration. If you assign an existing volume as the replication target, the data on that volume will be overwritten. The target volume can be greater or equal in size to the source volume, but it cannot be smaller.

Steps

1. Select **Management > Volumes**.
2. Click Actions icon for the volume you want to pair.
3. Click **Pair**.
4. In the **Pair Volume** dialog box, select **Start Pairing**.
5. Select **I Do Not** to indicate that you do not have access to the remote cluster.
6. Select a **Replication Mode** from the list:
 - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
 - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
 - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.
7. Click **Generate Key**.

Note: This action generates a text key for pairing and creates an unconfigured volume pair on the local cluster. If you do not complete the procedure, you will need to manually delete the volume pair.
8. Copy the pairing key to your computer's clipboard.
9. Make the pairing key accessible to the cluster admin at the remote cluster site.

Note: The volume pairing key should be treated in a secure manner and not used in a way that would allow accidental or unsecured access.



Attention: Do not modify any of the characters in the pairing key. The key becomes invalid if it is modified.

10. In the remote cluster Element UI, select **Management > Volumes**.
11. Click the Actions icon for the volume you want to pair.
12. Click **Pair**.
13. In the **Pair Volume** dialog box, select **Complete Pairing**.
14. Paste the pairing key from the other cluster into the **Pairing Key** box.
15. Click **Complete Pairing**.

After you confirm the pairing, the two clusters begin the process of connecting the volumes for pairing. During the pairing process, you can see messages in the **Volume Status** column of the **Volume Pairs** window. The volume pair displays `PausedMisconfigured` until the volume pair source and target are assigned.

Related tasks

[Assigning a replication source and target to paired volumes](#) on page 118

After volumes are paired, you must assign a source volume and its replication target volume. A replication source or target can be either volume in a volume pair. You can also use this procedure to redirect data sent to a source volume to a remote target volume should the source volume become unavailable.

Related reference

[Volume pairing messages](#) on page 122

You can view volume pairing messages during the initial pairing process from the Volume Pairs page under the Data Protection tab. These messages can display on both source and target ends of the pair in the Replicating Volumes list view.

[Volume pairing warnings](#) on page 123

The Volume Pairs page on the Data Protection tab provides these messages after you pair volumes. These messages can display on both source and target ends of the pair (unless otherwise indicated) in the Replicating Volumes list view.

Assigning a replication source and target to paired volumes

After volumes are paired, you must assign a source volume and its replication target volume. A replication source or target can be either volume in a volume pair. You can also use this procedure to redirect data sent to a source volume to a remote target volume should the source volume become unavailable.

Before you begin

You have access to the clusters containing the source and target volumes.

Steps

1. Prepare the source volume:
 - a. From the cluster that contains the volume you want to assign as source, select **Management > Volumes**.
 - b. Click the Actions icon for the volume you want to assign as source and click **Edit**.
 - c. In the **Access** drop-down list, select **Read/Write**.



Attention: If you are reversing source and target assignment, this action will cause the volume pair to display the following message until a new replication target is assigned:
`PausedMisconfigured`

Changing access pauses volume replication and causes the transmission of data to cease. Be sure that you have coordinated these changes at both sites.

- d. Click **Save Changes**.
2. Prepare the target volume:
 - a. From the cluster that contains the volume you want to assign as target, select **Management > Volumes**.
 - b. Click the Actions icon for the volume you want to assign as target and click **Edit**.
 - c. In the **Access** drop-down list, select **Replication Target**.



Attention: If you assign an existing volume as the replication target, the data on that volume will be overwritten. You should use a new target volume that contains no data and has the exact characteristics of the source volume, such as size, 512e setting, and QoS configuration. The target volume can be greater or equal in size to the source volume, but it cannot be smaller.

- d. Click **Save Changes**.

Related tasks

[Pairing volumes using a volume ID](#) on page 116

You can pair a volume with another volume on a remote cluster if you have cluster administrator credentials for the remote cluster.

[Pairing volumes using a pairing key](#) on page 117

If you do not have cluster admin credentials for a remote cluster, you can pair a volume with another volume on a remote cluster using a pairing key.

Validating volume replication

After a volume is replicated, you should ensure that the source and target volumes are active. When in an active state, volumes are paired, data is being sent from the source to the target volume, and the data is in sync.

Steps

1. From both clusters, select **Data Protection > Volume Pairs**.
2. Verify that the volume status is **Active**.

Related reference

[Volume pairing warnings](#) on page 123

The Volume Pairs page on the Data Protection tab provides these messages after you pair volumes. These messages can display on both source and target ends of the pair (unless otherwise indicated) in the Replicating Volumes list view.

Deleting a volume relationship after replication

After replication completes and you no longer need the volume pair relationship, you can delete the volume relationship.

Steps

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon for the volume pair you want to delete.

3. Click **Delete**.
4. Confirm the message.

Managing volume relationships

You can manage volume relationships in many ways, such as pausing replication, reversing volume pairing, changing the mode of replication, deleting a volume pair, or deleting a cluster pair.

Related tasks

[Pausing replication](#) on page 120

You can manually pause replication if you need to stop I/O processing for a short time. You might want to pause replication if there is a surge in I/O processing and you want to reduce the processing load.

[Changing the mode of replication](#) on page 120

You can edit volume pair properties to change the replication mode of the volume pair relationship.

[Deleting volume pairs](#) on page 121

You can delete a volume pair if want to remove a pair association between two volumes.

Pausing replication

You can manually pause replication if you need to stop I/O processing for a short time. You might want to pause replication if there is a surge in I/O processing and you want to reduce the processing load.

Steps

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon for the volume pair.
3. Click **Edit**.
4. In the **Edit Volume Pair** pane, manually pause the replication process.



Attention: Pausing or resuming volume replication manually causes the transmission of data to cease or resume. Be sure that you have coordinated these changes at both sites.

5. Click **Save Changes**.

Changing the mode of replication

You can edit volume pair properties to change the replication mode of the volume pair relationship.

Steps

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon for the volume pair.
3. Click **Edit**.
4. In the **Edit Volume Pair** pane, select a new replication mode:
 - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
 - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
 - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.



Attention: Changing the mode of replication changes the mode immediately. Be sure that you have coordinated these changes at both sites.

5. Click **Save Changes**.

Deleting volume pairs

You can delete a volume pair if want to remove a pair association between two volumes.

Steps

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon for the volume pair you want to delete.
3. Click **Delete**.
4. Confirm the message.

Deleting a cluster pair

You can delete a cluster pair from the Element UI of either of the clusters in the pair.

Steps

1. Click **Data Protection > Cluster Pairs**.
2. Click the Actions icon for a cluster pair.
3. In the resulting menu, click **Delete**.
4. Confirm the action.
5. Perform the steps again from the second cluster in the cluster pairing.

Cluster pair details

The Cluster Pairs page on the Data Protection tab provides information about clusters that have been paired or are in the process of being paired. The system displays pairing and progress messages in the Status column.

ID

A system-generated ID given to each cluster pair.

Remote Cluster Name

The name of the other cluster in the pair.

Remote MVIP

The management virtual IP address of the other cluster in the pair.

Status

Replication status of the remote cluster

Replicating Volumes

The number of volumes contained by the cluster that are paired for replication.

UUID

A unique ID given to each cluster in the pair.

Volume pair details

The Volume Pairs page on the Data Protection tab provides information about volumes that have been paired or are in the process of being paired. The system displays pairing and progress messages in the Volume Status column.

ID

System-generated ID for the volume.

Name

The name given to the volume when it was created. Volume names can be up to 223 characters and contain a-z, 0-9, and dash (-).

Account

Name of the account assigned to the volume.

Volume Status

Replication status of the volume

Snapshot Status

Status of the snapshot volume.

Mode

The client write replication method. Possible values are as follows:

- Async
- Snapshot-Only
- Sync

Direction

The direction of the volume data:

- Source volume icon () indicates data is being written to a target outside the cluster.
- Target volume icon () indicates data is being written to the local volume from an outside source.

Async Delay

Length of time since the volume was last synced with the remote cluster. If the volume is not paired, the value is null.

Remote Cluster

Name of the remote cluster on which the volume resides.

Remote Volume ID

Volume ID of the volume on the remote cluster.

Remote Volume Name

Name given to the remote volume when it was created.

Volume pairing messages

You can view volume pairing messages during the initial pairing process from the Volume Pairs page under the Data Protection tab. These messages can display on both source and target ends of the pair in the Replicating Volumes list view.

PausedDisconnected

Source replication or sync RPCs timed out. Connection to the remote cluster has been lost.
Check network connections to the cluster.

ResumingConnected

The remote replication sync is now active. Beginning the sync process and waiting for data.

ResumingRRSync

A single helix copy of the volume metadata is being made to the paired cluster.

ResumingLocalSync

A double helix copy of the volume metadata is being made to the paired cluster.

ResumingDataTransfer

Data transfer has resumed.

Active

Volumes are paired and data is being sent from the source to the target volume and the data is in sync.

Idle

No replication activity is occurring.

Volume pairing warnings

The Volume Pairs page on the Data Protection tab provides these messages after you pair volumes. These messages can display on both source and target ends of the pair (unless otherwise indicated) in the Replicating Volumes list view.

PausedClusterFull

Because the target cluster is full, source replication and bulk data transfer cannot proceed. The message displays on the source end of the pair only.

PausedExceededMaxSnapshotCount

The target volume already has the maximum number of snapshots and cannot replicate additional snapshots.

PausedManual

Local volume has been manually paused. It must be unpaused before replication resumes.

PausedManualRemote

Remote volume is in manual paused mode. Manual intervention required to unpause the remote volume before replication resumes.

PausedMisconfigured

Waiting for an active source and target. Manual intervention required to resume replication.

PausedQoS

Target QoS could not sustain incoming IO. Replication auto-resumes. The message displays on the source end of the pair only.

PausedSlowLink

Slow link detected and stopped replication. Replication auto-resumes. The message displays on the source end of the pair only.

PausedVolumeSizeMismatch

Target volume is not the same size as the source volume.

PausedXCopy

A SCSI XCOPY command is being issued to a source volume. The command must complete before replication can resume. The message displays on the source end of the pair only.

StoppedMisconfigured

A permanent configuration error has been detected. The remote volume has been purged or unpaired. No corrective action is possible; a new pairing must be established.

Using SnapMirror replication between Element and ONTAP clusters

You can create SnapMirror relationships from the Data Protection tab in the NetApp Element UI. SnapMirror functionality must be enabled to see this in the user interface.

About this task

IPv6 is not supported for SnapMirror replication between NetApp Element software and ONTAP clusters.

[NetApp video: SnapMirror for NetApp HCI and Element Software](#)

Systems running NetApp Element software support SnapMirror functionality to copy and restore Snapshot copies with NetApp ONTAP systems. The primary reason for using this technology is disaster recovery of NetApp HCI to ONTAP. Endpoints include ONTAP, ONTAP Select, and Cloud Volumes ONTAP. See TR-4641 NetApp HCI Data Protection.

[NetApp Technical Report 4641: NetApp HCI Data Protection](#)

Related information

[Building your Data Fabric with NetApp HCI, ONTAP, and Converged Infrastructure](#)

[Replication between NetApp Element Software and ONTAP](#)

SnapMirror overview

Systems running NetApp Element software support SnapMirror functionality to copy and restore snapshots with NetApp ONTAP systems.

Systems running Element can communicate directly with SnapMirror on ONTAP systems 9.3 or higher. The NetApp Element API provides methods to enable SnapMirror functionality on clusters, volumes, and snapshots. Additionally, the Element UI includes all necessary functionality to manage SnapMirror relationships between Element software and ONTAP systems.

You can replicate ONTAP originated volumes to Element volumes in specific use cases with limited functionality. For more information, see ONTAP documentation.

Related information

[Replication between Element software and ONTAP](#)

Enabling SnapMirror on the cluster

You must manually enable SnapMirror functionality at the cluster level through the NetApp Element UI. The system comes with SnapMirror functionality disabled by default, and it is not automatically enabled as part of a new installation or upgrade. Enabling the SnapMirror feature is a one-time configuration task.

Before you begin

The storage cluster must be running NetApp Element software.

About this task

SnapMirror can only be enabled for clusters running Element software used in conjunction with volumes on a NetApp ONTAP system. You should enable SnapMirror functionality only if your cluster is connected for use with NetApp ONTAP volumes.

Steps

1. Click **Clusters > Settings**.
2. Find the cluster-specific settings for SnapMirror.
3. Click **Enable SnapMirror**.

Note: Enabling SnapMirror functionality permanently changes the Element software configuration. You can disable the SnapMirror feature and restore the default settings only by returning the cluster to the factory image.

4. Click **Yes** to confirm the SnapMirror configuration change.

Enabling SnapMirror on the volume

You must enable SnapMirror on the volume in the Element UI. This allows replication of data to specified ONTAP volumes. This is permission from the administrator of the cluster running NetApp Element software for SnapMirror to control a volume.

Before you begin

- You have enabled SnapMirror in the Element UI for the cluster.
- A SnapMirror endpoint is available.
- The volume must be 512e block size.
- The volume is not participating in remote replication.
- The volume access type is not Replication Target.

Note: You can also set this property when creating or cloning a volume.

Steps

1. Click **Management > Volumes**.
2. Click the Actions icon for the volume you want to enable SnapMirror for.
3. In the resulting menu, select **Edit**.
4. In the **Edit Volume** dialog box, select the check box **Enable SnapMirror**.
5. Click **Save Changes**.

Creating an endpoint

You must create a SnapMirror endpoint in the NetApp Element UI before you can create a relationship.

Before you begin

- You have enabled SnapMirror in the Element UI for the storage cluster.
- You know the ONTAP credentials for the endpoint.

About this task

A SnapMirror endpoint is an ONTAP cluster that serves as a replication target for a cluster running Element software. Before you create a SnapMirror relationship, you first create a SnapMirror endpoint.

You can create and manage up to four SnapMirror endpoints on a storage cluster running Element software.

Note: If an existing endpoint was originally created using the API and credentials were not saved, you can see the endpoint in the Element UI and verify its existence, but it cannot be managed using the Element UI. This endpoint can then only be managed using the Element API. For information about the API methods, see API reference information.

[*Managing storage with the Element API*](#)

Steps

1. Click **Data Protection > SnapMirror Endpoints**.
2. Click **Create Endpoint**.

3. In the **Create a New Endpoint** dialog box, enter the cluster management IP address of the ONTAP system.
4. Enter the ONTAP administrator credentials associated with the endpoint.
5. Review additional details:
 - LIFs: Lists the ONTAP intercluster logical interfaces used to communicate with Element.
 - Status: Shows the current status of the SnapMirror endpoint. Possible values are: connected, disconnected, and unmanaged.
6. Click **Create Endpoint**.

Creating a SnapMirror relationship

You must create a SnapMirror relationship in the NetApp Element UI.

Before you begin

SnapMirror is enabled on the volume.

Note: When a volume is not yet enabled for SnapMirror and you select to create a relationship from the Element UI, SnapMirror is automatically enabled on that volume.

Steps

1. Click **Management > Volumes**.
2. Click the Actions icon for the volume that is to be a part of the relationship.
3. Click **Create a SnapMirror Relationship**.
4. In the **Create a SnapMirror Relationship** dialog box, select an endpoint from the **Endpoint** list.
5. Select if the relationship will be created using a new ONTAP volume or an existing ONTAP volume.
6. To create a new ONTAP volume in the Element UI, click **Create new volume**.
 - a. Select the **Storage Virtual Machine** for this relationship.
 - b. Select the **Aggregate** from the drop-down list.
 - c. In the **Volume Name Suffix** field, enter a suffix.

Note: The system detects the source volume name and copies it to the **Volume Name** field. The suffix you enter appends the name.

- d. Click **Create Destination Volume**.
7. To use an existing ONTAP volume, click **Use existing volume**.
 - a. Select the **Storage Virtual Machine** for this relationship.
 - b. Select the volume that is the destination for this new relationship.
8. In the **Relationship Details** section, select a policy. If the selected policy has keep rules, the Rules table displays the rules and associated labels.
9. Optional: Select a schedule.
This determines how often the relationship creates copies.
10. Optional: In the **Limit Bandwidth to** field, enter the maximum amount of bandwidth that can be consumed by data transfers associated with this relationship.
11. Review additional details:

State

Current relationship state of the destination volume. Possible values are:

- uninitialized: The destination volume has not been initialized.

- snapmirrored: The destination volume has been initialized and is ready to receive SnapMirror updates.
- broken-off: The destination volume is read/write and snapshots are present.

Status

Current status of the relationship. Possible values are idle, transferring, checking, quiescing, quiesced, queued, preparing, finalizing, aborting, and breaking.

Lag Time

The amount of time in seconds that the destination system lags behind the source system. The lag time must be no more than the transfer schedule interval.

Bandwidth Limit

The maximum amount of bandwidth that can be consumed by data transfers associated with this relationship.

Last Transferred

Timestamp of the last transferred snapshot. Click for further information.

Policy Name

The name of the ONTAP SnapMirror policy for the relationship.

Policy Type

Type of ONTAP SnapMirror policy selected for the relationship. Possible values are:

- async_mirror
- mirror_vault

Schedule Name

Name of the pre-existing schedule on the ONTAP system selected for this relationship.

12. To not initialize at this time, ensure that the **Initialize** check box is not selected.

Note: Initialization can be time-consuming. You might want to run this during off-peak hours. Initialization performs a baseline transfer; it makes a snapshot copy of the source volume, then transfers that copy and all the data blocks it references to the destination volume. You can initialize manually or use a schedule to start the initialization process (and subsequent updates) according to the schedule.

13. Click **Create Relationship**.
14. Click **Data Protection > SnapMirror Relationships** to view this new SnapMirror relationship.

SnapMirror relationship actions

You can configure a relationship from the SnapMirror Relationships page of the Data Protection tab. The options from the Actions icon are described here.

Edit

Edits the policy used or schedule for the relationship.

Delete

Deletes the SnapMirror relationship. This function does not delete the destination volume.

Initialize

Performs the first initial baseline transfer of data to establish a new relationship.

Update

Performs an on-demand update of the relationship, replicating any new data and Snapshot copies included since the last update to the destination.

Quiesce

Prevents any further updates for a relationship.

Resume

Resumes a relationship that is quiesced.

Break

Makes the destination volume read-write and stops all current and future transfers.
Determine that clients are not using the original source volume, because the reverse resync operation makes the original source volume read-only.

Resync

Reestablishes a broken relationship in the same direction before the break occurred.

Reverse Resync

Automates the necessary steps to create and initialize a new relationship in the opposite direction. This can be done only if the existing relationship is in a broken state. This operation will not delete the current relationship. The original source volume reverts to the most recent common Snapshot copy and resynchronizes with the destination. Any changes that are made to the original source volume since the last successful SnapMirror update are lost. Any changes that were made to, or new data written into the current destination volume is sent back to the original source volume.

Abort

Cancels a current transfer in progress. If a SnapMirror update is issued for an aborted relationship, the relationship continues with the last transfer from the last restart checkpoint that was created before the abort occurred.

SnapMirror labels

A SnapMirror label serves as a marker for transferring a specified snapshot according to the retention rules of the relationship.

Applying a label to a snapshot marks it as a target for SnapMirror replication. The role of the relationship is to enforce the rules upon data transfer by selecting the matching labeled snapshot, copying it to the destination volume, and ensuring the correct number of copies are kept. It refers to the policy to determine the keep count and the retention period. The policy can have any number of rules and each rule has a unique label. This label serves as the link between the snapshot and the retention rule.

It is the SnapMirror label that indicates which rule is applied for the selected snapshot, group snapshot, or schedule.

Adding SnapMirror labels to snapshots

SnapMirror labels specify the snapshot retention policy on the SnapMirror endpoint. You can add labels to snapshots and group snapshots.

Before you begin

- SnapMirror is enabled on the cluster.
- The label you want to add already exists in ONTAP.

About this task

You can view available labels from an existing SnapMirror relationship dialog box or the NetApp ONTAP System Manager.



Attention: When adding a label to a group snapshot, any existing labels to individual snapshots are overwritten.

Steps

1. Click **Data Protection > Snapshots or Group Snapshots** page.
2. Click the Actions icon for the snapshot or group snapshot you want to add a SnapMirror label to.
3. In the **Edit Snapshot** dialog box, enter text in the **SnapMirror Label** field. The label must match a rule label in the policy applied to the SnapMirror relationship.
4. Click **Save Changes**.

Adding SnapMirror labels to snapshot schedules

You can add SnapMirror labels to snapshot schedules to ensure that a SnapMirror policy is applied. You can view available labels from an existing SnapMirror relationship dialog box or the NetApp ONTAP System Manager.

Before you begin

- SnapMirror is enabled at the cluster level.
- The label you want to add already exists in ONTAP.

Steps

1. Click **Data Protection > Schedules**.
2. Add a SnapMirror label to a schedule in one of the following ways:

Option	Steps
Creating a new schedule	<ol style="list-style-type: none">a. Select Create Schedule.b. Enter all other relevant details.c. Select Create Schedule.
Modifying existing schedule	<ol style="list-style-type: none">a. Click the Actions icon for the schedule you want to add a label to and select Edit.b. In the resulting dialog box, enter text in the SnapMirror Label field.c. Select Save Changes.

Related tasks

[Creating a snapshot schedule](#) on page 108

You can schedule a snapshot of a volume or volumes to automatically occur at specified intervals.

Disaster recovery using SnapMirror

In the event of a problem with a volume or cluster running NetApp Element software, use the SnapMirror functionality to break the relationship and failover to the destination volume.

Note: If the original cluster has completely failed or is non-existent, contact NetApp Support for further assistance.

Performing a failover from an Element cluster

You can perform a failover from the Element cluster to make the destination volume read/write and accessible to hosts on the destination side. Before you perform a failover from the Element cluster, you must break the SnapMirror relationship.

Before you begin

- A SnapMirror relationship exists and has at least one valid snapshot on the destination volume.
- You have a need to failover to the destination volume due to unplanned outage or planned event at the primary site.

About this task

Use the NetApp Element UI to perform the failover. If the Element UI is not available, you can also use ONTAP System Manager or ONTAP CLI to issue the break relationship command.

Steps

1. In the Element UI, click **Data Protection > SnapMirror Relationships**.
2. Find the relationship with the source volume that you want to failover.
3. Click the Actions icon of this relationship.
4. Click **Break**.
5. Confirm the action.

The volume on the destination cluster now has read-write access and can be mounted to the application hosts to resume production workloads. All SnapMirror replication is halted as a result of this action. The relationship shows a state of broken-off.

Performing a failback to Element

When the issue on the primary side has been mitigated, you must resynchronize the original source volume and fail back to NetApp Element software. The steps you perform vary depending on whether the original source volume still exists or whether you need to failback to a newly created volume.

Related concepts

[SnapMirror failback scenarios](#) on page 130

The SnapMirror disaster recovery functionality is illustrated in two failback scenarios. These assume the original relationship has been failed over (broken).

Related tasks

[Performing a failback when source volume still exists](#) on page 132

You can resynchronize the original source volume and fail back using the NetApp Element UI. This procedure applies to scenarios where the original source volume still exists.

[Performing a failback when source volume no longer exists](#) on page 133

You can resynchronize the original source volume and fail back using the NetApp Element UI. This section applies to scenarios in which the original source volume has been lost but the original cluster is still intact. For instructions about how to restore to a new cluster, see the documentation on the NetApp Support Site.

SnapMirror failback scenarios

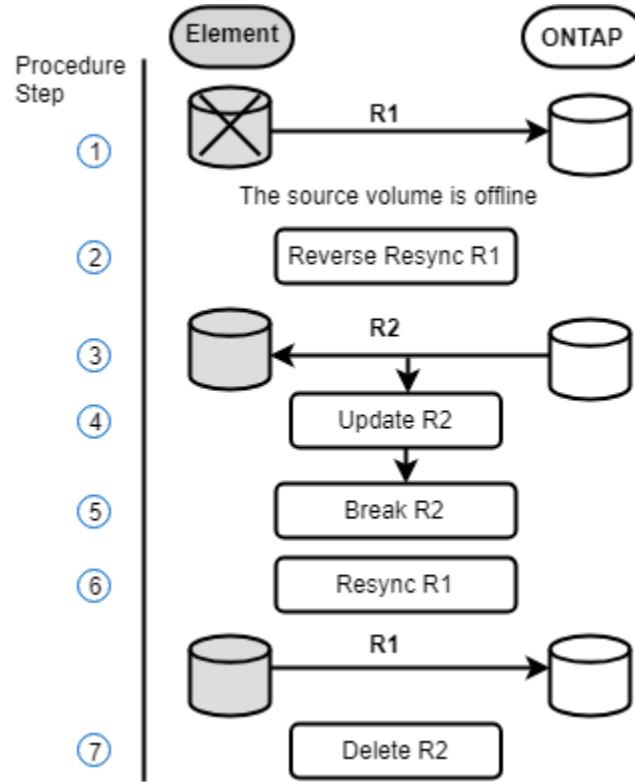
The SnapMirror disaster recovery functionality is illustrated in two failback scenarios. These assume the original relationship has been failed over (broken).

The steps from the corresponding procedures are added for reference.

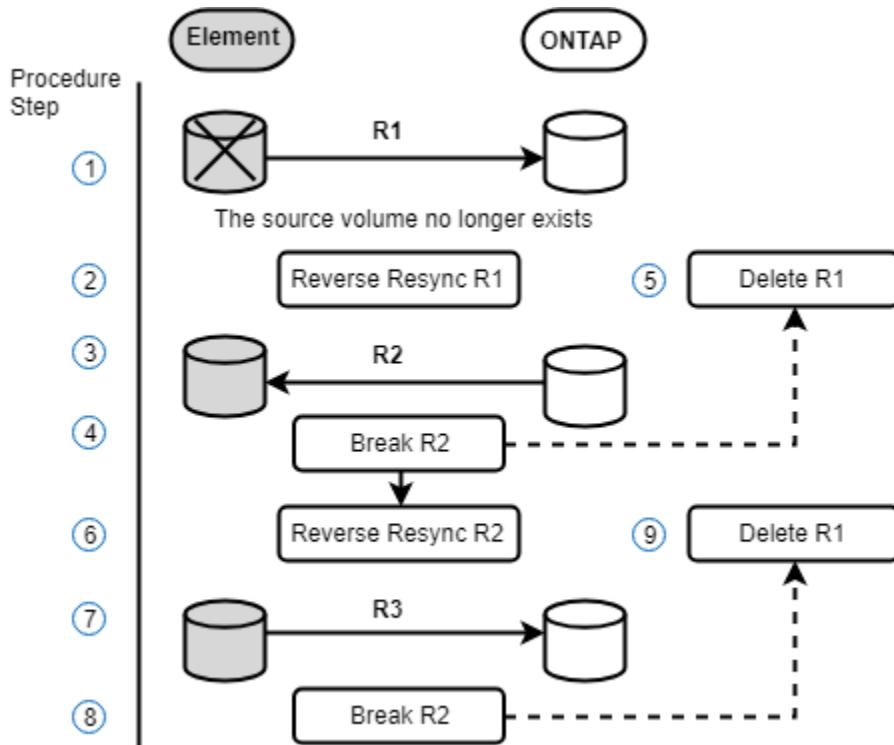
Note: In the examples here, R1 = the original relationship in which the cluster running NetApp Element software is the original source volume (Element) and ONTAP is the original

destination volume (ONTAP). R2 and R3 represent the inverse relationships created through the reverse resync operation.

The following image shows the failback scenario when the source volume still exists:



The following image shows the failback scenario when the source volume no longer exists:



Related tasks

[Performing a failback when source volume still exists](#) on page 132

You can resynchronize the original source volume and fail back using the NetApp Element UI. This procedure applies to scenarios where the original source volume still exists.

[Performing a failback when source volume no longer exists](#) on page 133

You can resynchronize the original source volume and fail back using the NetApp Element UI. This section applies to scenarios in which the original source volume has been lost but the original cluster is still intact. For instructions about how to restore to a new cluster, see the documentation on the NetApp Support Site.

Performing a failback when source volume still exists

You can resynchronize the original source volume and fail back using the NetApp Element UI. This procedure applies to scenarios where the original source volume still exists.

Steps

1. In the Element UI, find the relationship that you broke to perform the failover.
2. Click the Actions icon and click **Reverse Resync**.
3. Confirm the action.

Note: The Reverse Resync operation creates a new relationship in which the roles of the original source and destination volumes are reversed (this results in two relationships as the original relationship persists). Any new data from the original destination volume is transferred to the original source volume as part of the reverse resync operation. You can continue to access and write data to the active volume on the destination side, but you will need to disconnect all hosts to the source volume and perform a SnapMirror update before redirecting back to the original primary.

4. Click the Actions icon of the inverse relationship that you just created and click **Update**.

Now that you have completed the reverse resync and ensured that there are no active sessions connected to the volume on the destination side and that the latest data is on the original primary volume, you can perform the following steps to complete the failback and reactivate the original primary volume:

5. Click the Actions icon of the inverse relationship and click **Break**.
6. Click the Actions icon of the original relationship and click **Resync**.

Note: The original primary volume can now be mounted to resume production workloads on the original primary volume. The original SnapMirror replication resumes based on the policy and schedule configured for the relationship.

7. After you confirm that the original relationship status is "snapmirrored", click the Actions icon of the inverse relationship and click **Delete**.

Related concepts

[SnapMirror fallback scenarios](#) on page 130

The SnapMirror disaster recovery functionality is illustrated in two failback scenarios. These assume the original relationship has been failed over (broken).

Performing a failback when source volume no longer exists

You can resynchronize the original source volume and fail back using the NetApp Element UI. This section applies to scenarios in which the original source volume has been lost but the original cluster is still intact. For instructions about how to restore to a new cluster, see the documentation on the NetApp Support Site.

Before you begin

- You have a broken-off replication relationship between Element and ONTAP volumes.
- The Element volume is irretrievably lost.
- The original volume name shows as NOT FOUND.

Steps

1. In the Element UI, find the relationship that you broke to perform the failover.

Best Practices: Make note of the SnapMirror policy and schedule details of the original broken-off relationship. This information will be required when recreating the relationship.

2. Click the Actions icon and click **Reverse Resync**.
3. Confirm the action.

Note: The Reverse Resync operation creates a new relationship in which the roles of the original source volume and the destination volume are reversed (this results in two relationships as the original relationship persists). Because the original volume no longer exists, the system creates a new Element volume with the same volume name and volume size as the original source volume. The new volume is assigned a default QoS policy called sm-recovery and is associated with a default account called sm-recovery. You will want to manually edit the account and QoS policy for all volumes that are created by SnapMirror to replace the original source volumes that were destroyed.

Data from the latest snapshot is transferred to the new volume as part of the reverse resync operation. You can continue to access and write data to the active volume on the destination side, but you will need to disconnect all hosts to the active volume and perform a SnapMirror update before reinstating the original primary relationship in a later step. After you complete the reverse resync and ensure that there are no active sessions connected to the volume on the destination side and that the latest data is on the original primary volume, continue with the following steps to complete the failback and reactivate the original primary volume:

4. Click the Actions icon of the inverse relationship that was created during the Reverse Resync operation and click **Break**.
5. Click the Actions icon of the original relationship, in which the source volume does not exist, and click **Delete**.
6. Click the Actions icon of the inverse relationship, which you broke in step 4, and click **Reverse Resync**.
7. This reverses the source and destination and results in a relationship with the same volume source and volume destination as the original relationship.
8. Click the Actions icon and **Edit** to update this relationship with the original QoS policy and schedule settings you took note of.
9. Now it is safe to delete the inverse relationship that you reverse resynced in step 6.

Related concepts

[SnapMirror fallback scenarios](#) on page 130

The SnapMirror disaster recovery functionality is illustrated in two fallback scenarios. These assume the original relationship has been failed over (broken).

Performing a transfer or one-time migration from ONTAP to Element

Typically, when you use SnapMirror for disaster recovery from a SolidFire storage cluster running NetApp Element software to ONTAP software, Element is the source and ONTAP the destination. However, in some cases the ONTAP storage system can serve as the source and Element as the destination.

Before you begin

- Two scenarios exist:
 - No previous disaster recovery relationship exists. Follow all the steps in this procedure.
 - Previous disaster recovery relationship does exist, but not between the volumes being used for this mitigation. In this case, follow only steps 3 and 4 below.
- The Element destination node must have been made accessible to ONTAP.
- The Element volume must have been enabled for SnapMirror replication.

About this task

You must specify the Element destination path in the form `hostip:/lun/<id_number>`, where `lun` is the actual string “lun” and `id_number` is the ID of the Element volume.

Steps

1. Using ONTAP, create the relationship with the Element cluster:

```
snapmirror create -source-path SVM:volume|cluster://SVM/volume  
-destination-path hostip:/lun/name -type XDP -schedule schedule -policy  
policy
```

```
cluster_dst::> snapmirror create -source-path svm_1:volA_dst  
-destination-path 10.0.0.11:/lun/0005 -type XDP -schedule my_daily -policy MirrorLatest
```

2. Verify that the SnapMirror relationship was created by using the ONTAP `snapmirror show` command.

See information about creating a replication relationship in the ONTAP documentation and for complete command syntax, see the ONTAP man page.

3. Using the Element `CreateVolume` API, create the target volume and set the target volume access mode to SnapMirror:
Create an Element volume using the Element API

```
{  
    "method": "CreateVolume",  
    "params": {  
        "name": "SMTargetVolumeTest2",  
        "accountID": 1,  
        "totalSize": 100000000000,  
        "enable512e": true,  
        "attributes": {},  
        "qosPolicyID": 1,  
        "enableSnapMirrorReplication": true,  
        "access": "snapMirrorTarget"  
    },  
    "id": 1  
}
```

4. Initialize the replication relationship using the ONTAP `snapmirror initialize` command:

```
snapmirror initialize -source-path hostip:/lun/name  
-destination-path SVM:volume|cluster://SVM/volume
```

Backing up and restoring volumes

You can back up and restore volumes to other SolidFire storage, as well as secondary object stores that are compatible with Amazon S3 or OpenStack Swift.

About this task

When you restore volumes from OpenStack Swift or Amazon S3, you need manifest information from the original backup process. If you are restoring a volume that was backed up on a SolidFire storage system, no manifest information is required.

Related tasks

[Backing up a volume to an Amazon S3 object store](#) on page 135

You can back up volumes to external object stores that are compatible with Amazon S3.

[Backing up a volume to an OpenStack Swift object store](#) on page 136

You can back up volumes to external object stores that are compatible with OpenStack Swift.

[Backing up a volume to a SolidFire storage cluster](#) on page 136

You can back up volumes residing on a cluster to a remote cluster for storage clusters running Element software.

[Restoring a volume from backup on an Amazon S3 object store](#) on page 137

You can restore a volume from a backup on an Amazon S3 object store.

[Restoring a volume from backup on an OpenStack Swift object store](#) on page 137

You can restore a volume from a backup on an OpenStack Swift object store.

[Restoring a volume from backup on a SolidFire storage cluster](#) on page 138

You can restore a volume from a backup on a SolidFire storage cluster.

Backing up a volume to an Amazon S3 object store

You can back up volumes to external object stores that are compatible with Amazon S3.

Steps

1. Click **Management > Volumes**.
2. Click the Actions icon for the volume you want to back up.
3. In the resulting menu, click **Backup to**.
4. In the **Integrated Backup** dialog box under **Backup to**, select **S3**.
5. Select an option under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.

- **Uncompressed:** An uncompressed format compatible with other systems.

6. Enter a hostname to use to access the object store in the **Hostname** field.
7. Enter an access key ID for the account in the **Access Key ID** field.
8. Enter the secret access key for the account in the **Secret Access Key** field.
9. Enter the S3 bucket in which to store the backup in the **S3 Bucket** field.
10. Optional: Enter a nametag to append to the prefix in the **Nametag** field.
11. Click **Start Read**.

Backing up a volume to an OpenStack Swift object store

You can back up volumes to external object stores that are compatible with OpenStack Swift.

Steps

1. Click **Management > Volumes**.
2. Click the Actions icon for the volume to back up.
3. In the resulting menu, click **Backup to**.
4. In the **Integrated Backup** dialog box under **Backup to**, select **Swift**.
5. Select a data format under **Data Format**:
 - **Native:** A compressed format readable only by SolidFire storage systems.
 - **Uncompressed:** An uncompressed format compatible with other systems.
6. Enter a URL to use to access the object store in the **URL** field.
7. Enter a user name for the account in the **Username** field.
8. Enter the authentication key for the account in the **Authentication Key** field.
9. Enter the container in which to store the backup in the **Container** field.
10. Optional: Enter a name tag to append to the prefix in the **Nametag** field.
11. Click **Start Read**.

Backing up a volume to a SolidFire storage cluster

You can back up volumes residing on a cluster to a remote cluster for storage clusters running Element software.

Before you begin

Ensure that the source and target clusters are paired. See Cluster Pairing for more information.

About this task

When backing up or restoring from one cluster to another, the system generates a key to be used as authentication between the clusters. This bulk volume write key allows the source cluster to authenticate with the destination cluster, providing a level of security when writing to the destination volume. As part of the backup or restore process, you need to generate a bulk volume write key from the destination volume before starting the operation.

Steps

1. On the destination cluster, **Management > Volumes**.
2. Click the Actions icon for the destination volume.
3. In the resulting menu, click **Restore from**.
4. In the **Integrated Restore** dialog box, under **Restore from**, select **SolidFire**.
5. Select an option under **Data Format**:
 - **Native:** A compressed format readable only by SolidFire storage systems.

- **Uncompressed:** An uncompressed format compatible with other systems.

6. Click **Generate Key**.
7. Copy the key from the **Bulk Volume Write Key** box to your clipboard.
8. On the source cluster, go to **Management > Volumes**.
9. Click the Actions icon for the volume to back up.
10. In the resulting menu, click **Backup to**.
11. In the **Integrated Backup** dialog box under **Backup to**, select **SolidFire**.
12. Select the same option you selected earlier in the **Data Format** field.
13. Enter the management virtual IP address of the destination volume's cluster in the **Remote Cluster MVIP** field.
14. Enter the remote cluster user name in the **Remote Cluster Username** field.
15. Enter the remote cluster password in the **Remote Cluster Password** field.
16. In the **Bulk Volume Write Key** field, paste the key you generated on the destination cluster earlier.
17. Click **Start Read**.

Restoring a volume from backup on an Amazon S3 object store

You can restore a volume from a backup on an Amazon S3 object store.

Steps

1. Click **Reporting > Event Log**.
2. Locate the backup event that created the backup you need to restore.
3. In the **Details** column for the event, click **Show Details**.
4. Copy the manifest information to your clipboard.
5. Click **Management > Volumes**.
6. Click the Actions icon for the volume you want to restore.
7. In the resulting menu, click **Restore from**.
8. In the **Integrated Restore** dialog box under **Restore from**, select **S3**.
9. Select the option that matches the backup under **Data Format**:
 - **Native:** A compressed format readable only by SolidFire storage systems.
 - **Uncompressed:** An uncompressed format compatible with other systems.
10. Enter a hostname to use to access the object store in the **Hostname** field.
11. Enter an access key ID for the account in the **Access Key ID** field.
12. Enter the secret access key for the account in the **Secret Access Key** field.
13. Enter the S3 bucket in which to store the backup in the **S3 Bucket** field.
14. Paste the manifest information into the **Manifest** field.
15. Click **Start Write**.

Restoring a volume from backup on an OpenStack Swift object store

You can restore a volume from a backup on an OpenStack Swift object store.

Steps

1. Click **Reporting > Event Log**.
2. Locate the backup event that created the backup you need to restore.
3. In the **Details** column for the event, click **Show Details**.

4. Copy the manifest information to your clipboard.
5. Click **Management > Volumes**.
6. Click the Actions icon for the volume you want to restore.
7. In the resulting menu, click **Restore from**.
8. In the **Integrated Restore** dialog box under **Restore from**, select **Swift**.
9. Select the option that matches the backup under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
10. Enter a URL to use to access the object store in the **URL** field.
11. Enter a user name for the account in the **Username** field.
12. Enter the authentication key for the account in the **Authentication Key** field.
13. Enter the name of the container in which the backup is stored in the **Container** field.
14. Paste the manifest information into the **Manifest** field.
15. Click **Start Write**.

Restoring a volume from backup on a SolidFire storage cluster

You can restore a volume from a backup on a SolidFire storage cluster.

About this task

When backing up or restoring from one cluster to another, the system generates a key to be used as authentication between the clusters. This bulk volume write key allows the source cluster to authenticate with the destination cluster, providing a level of security when writing to the destination volume. As part of the backup or restore process, you need to generate a bulk volume write key from the destination volume before starting the operation.

Steps

1. On the destination cluster, click **Management > Volumes**.
2. Click the Actions icon for the volume you want to restore.
3. In the resulting menu, click **Restore from**.
4. In the **Integrated Restore** dialog box, under **Restore from**, select **SolidFire**.
5. Select the option that matches the backup under **Data Format**:
 - **Native**: A compressed format readable only by SolidFire storage systems.
 - **Uncompressed**: An uncompressed format compatible with other systems.
6. Click **Generate Key**.
7. Copy the **Bulk Volume Write Key** information to the clipboard.
8. On the source cluster, click **Management > Volumes**.
9. Click the Actions icon for the volume you want to use for the restore.
10. In the resulting menu, click **Backup to**.
11. In the **Integrated Backup** dialog box, select **SolidFire** under **Backup to**.
12. Select the option that matches the backup under **Data Format**.
13. Enter the management virtual IP address of the destination volume's cluster in the **Remote Cluster MVIP** field.
14. Enter the remote cluster user name in the **Remote Cluster Username** field.
15. Enter the remote cluster password in the **Remote Cluster Password** field.
16. Paste the key from your clipboard into the **Bulk Volume Write Key** field.

17. Click **Start Read**.

System monitoring and troubleshooting

You must monitor the system for diagnostic purposes and to get information about performance trends and statuses of various system operations. You might need to replace nodes or SSDs for maintenance purposes.

Related concepts

[Troubleshooting drives](#) on page 159

You can replace a failed solid-state drive (SSD) with a replacement drive. SSDs for SolidFire storage nodes are hot-swappable. If you suspect an SSD has failed, contact NetApp Support to verify the failure and walk you through the proper resolution procedure. NetApp Support also works with you to get a replacement drive according to your service-level agreement.

[Troubleshooting nodes](#) on page 162

You can remove nodes from a cluster for maintenance or replacement. You should use the NetApp Element UI or API to remove nodes before taking them offline.

[Working with per-node utilities for storage nodes](#) on page 163

You can use the per-node utilities to troubleshoot network problems if the standard monitoring tools in the NetApp Element software UI do not give you enough information for troubleshooting. Per-node utilities provide specific information and tools that can help you troubleshoot network problems between nodes or with the management node.

[Working with the management node](#) on page 169

You can use the management node (mNode) to upgrade system services, manage cluster assets and settings, run system tests and utilities, configure Active IQ for system monitoring, and enable NetApp Support access for troubleshooting.

[Understanding cluster fullness levels](#) on page 169

The cluster running Element software generates cluster faults to warn the storage administrator when the cluster is running out of capacity. There are three levels of cluster fullness, all of which are displayed in the NetApp Element UI: warning, error, and critical.

Related tasks

[Viewing information about system events](#) on page 141

You can view information about various events detected in the system. The system refreshes the event messages every 30 seconds. The event log displays key events for the cluster.

[Viewing system alerts](#) on page 144

You can view alerts for information about cluster faults or errors in the system. Alerts can be information, warnings, or errors and are a good indicator of how well the cluster is running. Most errors resolve themselves automatically.

[Viewing node performance activity](#) on page 155

You can view performance activity for each node in a graphical format. This information provides real-time statistics for CPU and read/write I/O operations per second (IOPS) for each drive the node. The utilization graph is updated every five seconds, and the drive statistics graph updates every ten seconds.

[Viewing volume performance](#) on page 156

You can view detailed performance information for all volumes in the cluster. You can sort the information by volume ID or by any of the performance columns. You can also use filter the information by certain criteria.

[Viewing iSCSI sessions](#) on page 157

You can view the iSCSI sessions that are connected to the cluster. You can filter the information to include only the desired sessions.

[Viewing Fibre Channel sessions](#) on page 158

You can view the Fibre Channel (FC) sessions that are connected to the cluster. You can filter information to include only those connections you want displayed in the window.

[Enabling FIPS 140-2 for HTTPS on your cluster](#) on page 60

You can use the `EnableFeature` API method to enable the FIPS 140-2 operating mode for HTTPS communications.

Related reference

[Viewing status of running tasks](#) on page 143

You can view the progress and completion status of running tasks in the web UI that are being reported by the `ListSyncJobs` and `ListBulkVolumeJobs` API methods. You can access the Running Tasks page from the Reporting tab of the Element UI.

Viewing information about system events

You can view information about various events detected in the system. The system refreshes the event messages every 30 seconds. The event log displays key events for the cluster.

Step

In the Element UI, select **Reporting > Event Log**.

For every event, you see the following information:

Item	Description
ID	Unique ID associated with each event.
Event Type	The type of event being logged, for example, API events or clone events.
Message	Message associated with the event.
Details	Information that helps identify why the event occurred.
Service ID	The service that reported the event (if applicable).
Node	The node that reported the event (if applicable).
Drive ID	The drive that reported the event (if applicable).
Event Time	The time the event occurred.

Related reference

[Event types](#) on page 141

The system reports multiple types of events; each event is an operation that the system has completed. Events can be routine, normal events or events that require administrator attention. The Event Types column on the Event Log page indicates in which part of the system the event occurred.

Event types

The system reports multiple types of events; each event is an operation that the system has completed. Events can be routine, normal events or events that require administrator attention. The Event Types column on the Event Log page indicates in which part of the system the event occurred.

Note: The system does not log read-only API commands in the event log.

The following list describes the types of events that appear in the event log:

apiEvent

Events initiated by a user through an API or web UI that modify settings.

binAssignmentsEvent

Events related to the assignment of data bins. Bins are essentially containers that hold data and are mapped across the cluster.

binSyncEvent

System events related to a reassignment of data among block services.

bsCheckEvent

System events related to block service checks.

bsKillEvent

System events related to block service terminations.

bulkOpEvent

Events related to operations performed on an entire volume, such as a backup, restore, snapshot, or clone.

cloneEvent

Events related to volume cloning.

clusterMasterEvent

Events appearing upon cluster initialization or upon configuration changes to the cluster, such as adding or removing nodes.

csumEvent

Events related to invalid data checksums on the disk.

dataEvent

Events related to reading and writing data.

dbEvent

Events related to the global database maintained by ensemble nodes in the cluster.

driveEvent

Events related to drive operations.

encryptionAtRestEvent

Events related to the process of encryption on a cluster.

ensembleEvent

Events related to increasing or decreasing the number of nodes in an ensemble.

fibreChannelEvent

Events related to the configuration of and connections to the nodes.

gcEvent

Events related to processes run every 60 minutes to reclaim storage on block drives. This process is also known as garbage collection.

ieEvent

Internal system error.

installEvent

Automatic software installation events. Software is being automatically installed on a pending node.

iSCSIEvent

Events related to iSCSI issues in the system.

limitEvent

Events related to the number of volumes or virtual volumes in an account or in the cluster nearing the maximum allowed.

networkEvent

Events related to the status of virtual networking.

platformHardwareEvent

Events related to issues detected on hardware devices.

remoteClusterEvent

Events related to remote cluster pairing.

schedulerEvent

Events related to scheduled snapshots.

serviceEvent

Events related to system service status.

sliceEvent

Events related to the Slice Server, such as removing a metadata drive or volume.

snmpTrapEvent

Events related to SNMP traps.

statEvent

Events related to system statistics.

tsEvent

Events related to the system transport service.

unexpectedException

Events related to unexpected system exceptions.

ureEvent

Events related to Unrecoverable Read Errors that occur while reading from the storage device.

vasaProviderEvent

Events related to a VASA (vSphere APIs for Storage Awareness) Provider.

Viewing status of running tasks

You can view the progress and completion status of running tasks in the web UI that are being reported by the `ListSyncJobs` and `ListBulkVolumeJobs` API methods. You can access the Running Tasks page from the Reporting tab of the Element UI.

If there are a large number of tasks, the system might queue them and run them in batches. The Running Tasks page displays the services currently being synchronized. When a task is complete, it is replaced by the next queued synchronizing task. Synchronizing tasks might continue to appear on the Running Tasks page until there are no more tasks to complete.

Note: You can see replication synchronizations data for volumes undergoing replication on the Running Tasks page of the cluster containing the target volume.

Viewing system alerts

You can view alerts for information about cluster faults or errors in the system. Alerts can be information, warnings, or errors and are a good indicator of how well the cluster is running. Most errors resolve themselves automatically.

About this task

You can use the `ListClusterFaults` API method to automate alert monitoring. This enables you to be notified about all alerts that occur.

Steps

1. In the Element UI, select **Reporting > Alerts**.

The system refreshes the alerts on the page every 30 seconds.

For every event, you see the following information:

Item	Description
ID	Unique ID associated with a cluster alert.
Severity	The degree of importance of the alert. Possible values: <ul style="list-style-type: none">• warning: A minor issue that might soon require attention. System upgrades are still allowed.• error: A failure that might cause performance degradation or loss of high availability (HA). Errors generally should not affect service otherwise.• critical: A serious failure that affects service. The system is unable to serve API or client I/O requests. Operating in this state could lead to potential loss of data.• bestPractice: A recommended system configuration best practice is not being used.
Type	The component that the fault affects. Can be node, drive, cluster, service, or volume.
Node	Node ID for the node that this fault refers to. Included for node and drive faults, otherwise set to - (dash).
Drive ID	Drive ID for the drive that this fault refers to. Included for drive faults, otherwise set to - (dash).
Error Code	A descriptive code that indicates what caused the fault.
Details	A description of the fault with additional details.
Date	The date and time the fault was logged.

2. Click **Show Details** for an individual alert to view information about the alert.

3. To view the details of all alerts on the page, click the Details column.

After the system resolves an alert, all information about the alert including the date it was resolved is moved to the Resolved area.

Related reference

[Cluster fault codes](#) on page 145

The system reports an error or a state that might be of interest by generating a fault code, which is listed on the Alerts page. These codes help you determine what component of the system experienced the alert and why the alert was generated.

Related information

[Managing storage with the Element API](#)

Cluster fault codes

The system reports an error or a state that might be of interest by generating a fault code, which is listed on the Alerts page. These codes help you determine what component of the system experienced the alert and why the alert was generated.

The following list outlines the different types of codes:

authenticationServiceFault

The Authentication Service on one or more cluster nodes is not functioning as expected.

Contact NetApp Support for assistance.

availableVirtualNetworkIPAddressesLow

The number of virtual network addresses in the block of IP addresses is low.

To resolve this fault, add more IP addresses to the block of virtual network addresses.

blockClusterFull

There is not enough free block storage space to support a single node loss. See the `GetClusterFullThreshold` API method for details on cluster fullness levels. This cluster fault indicates one of the following conditions:

- `stage3Low` (Warning): User-defined threshold was crossed. Adjust Cluster Full settings or add more nodes.
- `stage4Critical` (Error): There is not enough space to recover from a 1-node failure. Creation of volumes, snapshots, and clones is not allowed.
- `stage5CompletelyConsumed` (Critical): No writes or new iSCSI connections are allowed. Current iSCSI connections will be maintained. Writes will fail until more capacity is added to the cluster.

To resolve this fault, purge or delete volumes or add another storage node to the storage cluster.

blocksDegraded

Block data is no longer fully replicated due to a failure.

Severity	Description
Error	Only a single copy of the complete block data is still available.
Critical	No complete copy of the block data is available.

To resolve this fault, restore any offline nodes or block services, or contact NetApp Support.

blockServiceTooFull

A block service is using too much space.

To resolve this fault, add more provisioned capacity.

blockServiceUnhealthy

A block service has been detected as unhealthy:

- Severity = Warning: No action is taken. This warning period will expire in `cTimeUntilBSIsKilledMSec=330000` milliseconds.
- Severity = Error: The system is automatically decommissioning data and re-replicating its data to other healthy drives.
- Severity = Critical: There are failed block services on several nodes greater than or equal to the replication count (2 for double helix). Data is unavailable and bin syncing will not finish.

Check for network connectivity issues and hardware errors. There will be other faults if specific hardware components have failed. The fault will clear when the block service is accessible or when the service has been decommissioned.

clockSkewExceedsFaultThreshold

Time skew between the Cluster master and the node which is presenting a token exceeds the recommended threshold. Storage cluster cannot correct the time skew between the nodes automatically.

To resolve this fault, use NTP servers that are internal to your network, rather than the installation defaults. If you are using an internal NTP server, contact NetApp Support for assistance.

clusterCannotSync

There is an out-of-space condition and data on the offline block storage drives cannot be synced to drives that are still active.

To resolve this fault, add more storage.

clusterFull

There is no more free storage space in the storage cluster.

To resolve this fault, add more storage.

clusterIOPSAreOverProvisioned

Cluster IOPS are over provisioned. The sum of all minimum QoS IOPS is greater than the expected IOPS of the cluster. Minimum QoS cannot be maintained for all volumes simultaneously.

To resolve this issue, lower the minimum QoS IOPS settings for volumes.

disableDriveSecurityFailed

The cluster is not configured to enable drive security (Encryption at Rest), but at least one drive has drive security enabled, meaning that disabling drive security on those drives failed. This fault is logged with “Warning” severity.

To resolve this fault, check the fault details for the reason why drive security could not be disabled. Possible reasons are:

- The encryption key could not be acquired, investigate the problem with access to the key or the external key server.
- The disable operation failed on the drive, determine whether the wrong key could possibly have been acquired.

If neither of these are the reason for the fault, the drive might need to be replaced.

You can attempt to recover a drive that does not successfully disable security even when the correct authentication key is provided. To perform this operation, remove the drive(s) from the system by moving it to Available, perform a secure erase on the drive and move it back to Active.

disconnectedClusterPair

A cluster pair is disconnected or configured incorrectly. Check network connectivity between the clusters.

disconnectedRemoteNode

A remote node is either disconnected or configured incorrectly. Check network connectivity between the nodes.

disconnectedSnapMirrorEndpoint

A remote SnapMirror endpoint is disconnected or configured incorrectly. Check network connectivity between the cluster and the remote SnapMirrorEndpoint.

driveAvailable

One or more drives are available in the cluster. In general, all clusters should have all drives added and none in the available state. If this fault appears unexpectedly, contact NetApp Support.

To resolve this fault, add any available drives to the storage cluster.

driveFailed

The cluster returns this fault when one or more drives have failed, indicating one of the following conditions:

- The drive manager cannot access the drive.
- The slice or block service has failed too many times, presumably because of drive read or write failures, and cannot restart.
- The drive is missing.
- The master service for the node is inaccessible (all drives in the node are considered missing/failed).
- The drive is locked and the authentication key for the drive cannot be acquired.
- The drive is locked and the unlock operation fails.

To resolve this issue:

- Check network connectivity for the node.
- Replace the drive.
- Ensure that the authentication key is available.

driveWearFault

A drive's remaining life has dropped below thresholds, but it is still functioning. There are two possible severity levels for this fault: Critical and Warning:

- Drive with serial: <serial number> in slot: <node slot><drive slot> has critical wear levels.
- Drive with serial: <serial number> in slot: <node slot><drive slot> has low wear reserves.

To resolve this fault, replace the drive soon.

duplicateClusterMasterCandidates

More than one storage cluster master candidate has been detected. Contact NetApp Support for assistance.

enableDriveSecurityFailed

The cluster is configured to require drive security (Encryption at Rest), but drive security could not be enabled on at least one drive. This fault is logged with “Warning” severity.

To resolve this fault, check the fault details for the reason why drive security could not be enabled. Possible reasons are:

- The encryption key could not be acquired, investigate the problem with access to the key or the external key server.
- The enable operation failed on the drive, determine whether the wrong key could possibly have been acquired.

If neither of these are the reason for the fault, the drive might need to be replaced.

You can attempt to recover a drive that does not successfully enable security even when the correct authentication key is provided. To perform this operation, remove the drive(s) from the system by moving it to Available, perform a secure erase on the drive and move it back to Active.

ensembleDegraded

Network connectivity or power has been lost to one or more of the ensemble nodes.

To resolve this fault, restore network connectivity or power.

exception

A fault reported that is other than a routine fault. These faults are not automatically cleared from the fault queue. Contact NetApp Support for assistance.

failedSpaceTooFull

A block service is not responding to data write requests. This causes the slice service to run out of space to store failed writes.

To resolve this fault, restore block services functionality to allow writes to continue normally and failed space to be flushed from the slice service.

fanSensor

A fan sensor has failed or is missing.

To resolve this fault, replace any failed hardware.

fibreChannelAccessDegraded

A Fibre Channel node is not responding to other nodes in the storage cluster over its storage IP for a period of time. In this state, the node will then be considered unresponsive and generate a cluster fault. Check network connectivity.

fibreChannelAccessUnavailable

All Fibre Channel nodes are unresponsive. The node IDs are displayed. Check network connectivity.

fibreChannelActiveIxL

The IxL Nexus count is approaching the supported limit of 8000 active sessions per Fibre Channel node.

- Best practice limit is 5500.
- Warning limit is 7500.
- Maximum limit (not enforced) is 8192.

To resolve this fault, reduce the IxL Nexus count below the best practice limit of 5500.

fibreChannelConfig

This cluster fault indicates one of the following conditions:

- There is an unexpected Fibre Channel port on a PCI slot.
- There is an unexpected Fibre Channel HBA model.
- There is a problem with the firmware of a Fibre Channel HBA.
- A Fibre Channel port is not online.
- There is a persistent issue configuring Fibre Channel passthrough.

Contact NetApp Support for assistance.

fibreChannelStaticIxL

The IxL Nexus count is approaching the supported limit of 16000 static sessions per Fibre Channel node.

- Best practice limit is 11000.

- Warning limit is 15000.
- Maximum limit (enforced) is 16384.

To resolve this fault, reduce the IxL Nexus count below the best practice limit of 11000.

fileSystemCapacityLow

There is insufficient space on one of the filesystems.

To resolve this fault, add more capacity to the filesystem.

fipsDrivesMismatched

A non-FIPS drive has been physically inserted into a FIPS capable storage node or a FIPS drive has been physically inserted into a non-FIPS storage node. A single fault is generated per node and lists all drives affected.

To resolve this fault, remove or replace the mismatched drive or drives in question.

fipsDrivesOutOfCompliance

The system has detected that Encryption at Rest was disabled after the FIPS Drives feature was enabled. This fault is also generated when the FIPS Drives feature is enabled and a non-FIPS drive or node is present in the storage cluster.

To resolve this fault, enable Encryption at Rest or remove the non-FIPS hardware from the storage cluster.

fipsSelfTestFailure

The FIPS subsystem has detected a failure during the self test.

Contact NetApp Support for assistance.

hardwareConfigMismatch

This cluster fault indicates one of the following conditions:

- The configuration does not match the node definition.
- There is an incorrect drive size for this type of node.
- An unsupported drive has been detected. A possible reason is that the installed Element version does not recognize this drive. Recommend updating the Element software on this node.
- There is a drive firmware mismatch.
- The drive encryption capable state does not match the node.

Contact NetApp Support for assistance.

idPCertificateExpiration

The cluster's service provider SSL certificate for use with a third-party identity provider (IdP) is nearing expiration or has already expired. This fault uses the following severities based on urgency:

Severity	Description
Warning	Certificate expires within 30 days.
Error	Certificate expires within 7 days.
Critical	Certificate expires within 3 days or has already expired.

To resolve this fault, update the SSL certificate before it expires. Use the `UpdateIdpConfiguration` API method with `refreshCertificateExpirationTime=true` to provide the updated SSL certificate.

inconsistentBondModes

The bond modes on the VLAN device are missing. This fault will display the expected bond mode and the bond mode currently in use.

inconsistentInterfaceConfiguration

The interface configuration is inconsistent.

To resolve this fault, ensure the node interfaces in the storage cluster are consistently configured.

inconsistentMtus

This cluster fault indicates one of the following conditions:

- Bond1G mismatch: Inconsistent MTUs have been detected on Bond1G interfaces.
- Bond10G mismatch: Inconsistent MTUs have been detected on Bond10G interfaces.

This fault displays the node or nodes in question along with the associated MTU value.

inconsistentRoutingRules

The routing rules for this interface are inconsistent.

inconsistentSubnetMasks

The network mask on the VLAN device does not match the internally recorded network mask for the VLAN. This fault displays the expected network mask and the network mask currently in use.

incorrectBondPortCount

The number of bond ports is incorrect.

invalidConfiguredFibreChannelNodeCount

One of the two expected Fibre Channel node connections is degraded. This fault appears when only one Fibre Channel node is connected.

To resolve this fault, check the cluster network connectivity and network cabling, and check for failed services. If there are no network or service problems, contact NetApp Support for a Fibre Channel node replacement.

irqBalanceFailed

An exception occurred while attempting to balance interrupts.

Contact NetApp Support for assistance.

kmipCertificateFault

- Root Certification Authority (CA) certificate is nearing expiration.

To resolve this fault, acquire a new certificate from the root CA with expiration date at least 30 days out and use `ModifyKeyServerKmip` to provide the updated root CA certificate.

- Client certificate is nearing expiration.

To resolve this fault, create a new CSR using `GetClientCertificateSigningRequest`, have it signed ensuring the new expiration date is at least 30 days out, and use `ModifyKeyServerKmip` to replace the expiring KMIP client certificate with the new certificate.

- Root Certification Authority (CA) certificate has expired.

To resolve this fault, acquire a new certificate from the root CA with expiration date at least 30 days out and use `ModifyKeyServerKmip` to provide the updated root CA certificate.

- Client certificate has expired.

To resolve this fault, create a new CSR using `GetClientCertificateSigningRequest`, have it signed ensuring the new expiration date is at least 30 days out, and use `ModifyKeyServerKmip` to replace the expired KMIP client certificate with the new certificate.

- Root Certification Authority (CA) certificate error.

To resolve this fault, check that the correct certificate was provided, and, if needed, reacquire the certificate from the root CA. Use `ModifyKeyServerKmip` to install the correct KMIP client certificate.

- Client certificate error.

To resolve this fault, check that the correct KMIP client certificate is installed. The root CA of the client certificate should be installed on the EKS. Use `ModifyKeyServerKmip` to install the correct KMIP client certificate.

kmipServerFault

- Connection failure

To resolve this fault, check that the External Key Server is alive and reachable via the network. Use `TestKeyServerKmip` and `TestKeyProviderKmip` to test your connection.

- Authentication failure

To resolve this fault, check that the correct root CA and KMIP client certificates are being used, and that the private key and the KMIP client certificate match.

- Server error

To resolve this fault, check the details for the error. Troubleshooting on the External Key Server might be necessary based on the error returned.

memoryEccThreshold

A large number of correctable or uncorrectable ECC errors have been detected. When a severity of type `Error` is returned, this is likely due to a DIMM failure.

Contact NetApp Support for assistance.

memoryUsageThreshold

Memory usage is above normal.

Contact NetApp Support for assistance.

metadataClusterFull

There is not enough free metadata storage space to support a single node loss. See the `GetClusterFullThreshold` API method for details on cluster fullness levels. This cluster fault indicates one of the following conditions:

- `stage3Low` (Warning): User-defined threshold was crossed. Adjust Cluster Full settings or add more nodes.
- `stage4Critical` (Error): There is not enough space to recover from a 1-node failure. Creation of volumes, snapshots, and clones is not allowed.
- `stage5CompletelyConsumed` (Critical): No writes or new iSCSI connections are allowed. Current iSCSI connections will be maintained. Writes will fail until more capacity is added to the cluster. Purge or delete data or add more nodes.

To resolve this fault, purge or delete volumes or add another storage node to the storage cluster.

mtuCheckFailure

A network device is not configured for the proper MTU size.

To resolve this fault, ensure that all network interfaces and switch ports are configured for jumbo frames (MTUs up to 9000 bytes in size).

networkConfig

This cluster fault indicates one of the following conditions:

- An expected interface is not present.
- A duplicate interface is present.
- A configured interface is down.
- A network restart is required.

Contact NetApp Support for assistance.

noAvailableVirtualNetworkIPAddresses

There are no available virtual network addresses in the block of IP addresses. No more storage nodes can be added to the cluster.

To resolve this fault, add more IP addresses to the block of virtual network addresses.

nodeHardwareFault (Network interface <name> is down or cable is unplugged)

A network interface is either down or the cable is unplugged.

To resolve this fault, check network connectivity for the node or nodes.

nodeHardwareFault (Drive encryption capable state mismatches node's encryption capable state for the drive in slot <node slot><drive slot>)

A drive does not match encryption capabilities with the storage node it is installed in.

nodeHardwareFault (Incorrect <drive type> drive size <actual size> for the drive in slot <node slot><drive slot> for this node type - expected <expected size>)

A storage node contains a drive that is the incorrect size for this node.

nodeHardwareFault (Unsupported drive detected in slot <node slot><drive slot>; drive statistics and health information will be unavailable)

A storage node contains a drive it does not support.

nodeHardwareFault (The drive in slot <node slot><drive slot> should be using firmware version <expected version>, but is using unsupported version <actual version>)

A storage node contains a drive running an unsupported firmware version.

nodeOffline

Element software cannot communicate with the specified node. Check network connectivity.

notUsingLACPBondMode

LACP bonding mode is not configured.

To resolve this fault, use LACP bonding when deploying storage nodes; clients might experience performance issues if LACP is not enabled and properly configured.

ntpServerUnreachable

The storage cluster cannot communicate with the specified NTP server or servers.

To resolve this fault, check the configuration for the NTP server, network, and firewall.

ntpTimeNotInSync

The difference between storage cluster time and the specified NTP server time is too large.

The storage cluster cannot correct the difference automatically.

To resolve this fault, use NTP servers that are internal to your network, rather than the installation defaults. If you are using internal NTP servers and the issue persists, contact NetApp Support for assistance.

nvrampDeviceStatus

An NVRAM device has an error, is failing, or has failed. This fault uses the following severities based on urgency:

Severity	Description
Warning	A warning has been detected by the hardware. This condition may be transitory such as an over temperature warning.
Error	An Error status has been detected by the hardware. The Cluster Master attempts to remove the slice drive from operation. If secondary slice services are not available, the drive will not be removed.
Critical	A Critical status has been detected by the hardware. The Cluster Master attempts to remove the slice drive from operation. If secondary slice services are not available, the drive will not be removed.

To resolve this issue, replace any failed hardware.

powerSupplyError

This cluster fault indicates one of the following conditions:

- A power supply is not present.
- A power supply has failed.
- A power supply input is missing or out of range.

To resolve this fault, verify that redundant power is supplied to all nodes. Contact NetApp Support for assistance.

provisionedSpaceTooFull

The overall provisioned capacity of the cluster is too full.

To resolve this fault, add more provisioned space, or delete and purge volumes.

remoteRepAsyncDelayExceeded

The configured asynchronous delay for replication has been exceeded. Check network connectivity between clusters.

remoteRepClusterFull

The volumes have paused remote replication because the target storage cluster is too full.

To resolve this fault, free up some space on the target storage cluster.

remoteRepSnapshotClusterFull

The volumes have paused remote replication of snapshots because the target storage cluster is too full.

To resolve this fault, free up some space on the target storage cluster.

remoteRepSnapshotsExceededLimit

The volumes have paused remote replication of snapshots because the target storage cluster volume has exceeded its snapshot limit.

To resolve this fault, increase the snapshot limit on the target storage cluster.

scheduleActionError

One or more of the scheduled activities ran, but failed.

The fault clears if the scheduled activity runs again and succeeds, if the scheduled activity is deleted, or if the activity is paused and resumed.

sensorReadingFailed

The Baseboard Management Controller (BMC) self-test failed or a sensor could not communicate with the BMC.

Contact NetApp Support for assistance.

serviceNotRunning

A required service is not running.

Contact NetApp Support for assistance.

sliceServiceTooFull

A slice service has too little provisioned capacity assigned to it.

To resolve this fault, add more provisioned capacity.

sliceServiceUnhealthy

The system has detected that a slice service is unhealthy and is automatically decommissioning it.

- Severity = Warning: No action is taken. This warning period will expire in 6 minutes.
- Severity = Error: The system is automatically decommissioning data and re-replicating its data to other healthy drives.

Check for network connectivity issues and hardware errors. There will be other faults if specific hardware components have failed. The fault will clear when the slice service is accessible or when the service has been decommissioned.

sshEnabled

The SSH service is enabled on one or more nodes in the storage cluster.

To resolve this fault, disable the SSH service on the appropriate node or nodes or contact NetApp Support for assistance.

sslCertificateExpiration

The SSL certificate associated with this node is nearing expiration or has expired. This fault uses the following severities based on urgency:

Severity	Description
Warning	Certificate expires within 30 days.
Error	Certificate expires within 7 days.
Critical	Certificate expires within 3 days or has already expired.

To resolve this fault, renew the SSL certificate. If needed, contact NetApp Support for assistance.

strandedCapacity

A single node accounts for more than half of the storage cluster capacity.

In order to maintain data redundancy, the system reduces the capacity of the largest node so that some of its block capacity is stranded (not used).

To resolve this fault, add more drives to existing storage nodes or add storage nodes to the cluster.

tempSensor

A temperature sensor is reporting higher than normal temperatures. This fault can be triggered in conjunction with powerSupplyError or fanSensor faults.

To resolve this fault, check for airflow obstructions near the storage cluster. If needed, contact NetApp Support for assistance.

upgrade

An upgrade has been in progress for more than 24 hours.

To resolve this fault, resume the upgrade or contact NetApp Support for assistance.

unresponsiveService

A service has become unresponsive.

Contact NetApp Support for assistance.

virtualNetworkConfig

This cluster fault indicates one of the following conditions:

- An interface is not present.
- There is an incorrect namespace on an interface.
- There is an incorrect netmask.
- There is an incorrect IP address.
- An interface is not up and running.
- There is a superfluous interface on a node.

Contact NetApp Support for assistance.

volumeDegraded

Secondary volumes have not finished replicating and synchronizing. The message is cleared when the synchronizing is complete.

volumesOffline

One or more volumes in the storage cluster are offline. The **volumeDegraded** fault will also be present.

Contact NetApp Support for assistance.

Viewing node performance activity

You can view performance activity for each node in a graphical format. This information provides real-time statistics for CPU and read/write I/O operations per second (IOPS) for each drive the node. The utilization graph is updated every five seconds, and the drive statistics graph updates every ten seconds.

Steps

1. Click **Cluster > Nodes**.
2. Click **Actions** for the node you want to view.
3. Click **View Details**.

Note: You can see specific points in time on the line and bar graphs by positioning your cursor over the line or bar.

Viewing volume performance

You can view detailed performance information for all volumes in the cluster. You can sort the information by volume ID or by any of the performance columns. You can also use filter the information by certain criteria.

About this task

You can change how often the system refreshes performance information on the page by clicking the **Refresh every** list, and choosing a different value. The default refresh interval is 10 seconds if the cluster has less than 1000 volumes; otherwise, the default is 60 seconds. If you choose a value of Never, automatic page refreshing is disabled.

You can reenable automatic refreshing by clicking **Turn on auto-refresh**.

Steps

1. In the Element UI, select **Reporting > Volume Performance**.
2. In the volume list, click the Actions icon for a volume.
3. Click **View Details**.
A tray is displayed at the bottom of the page containing general information about the volume.
4. To see more detailed information about the volume, click **See More Details**.
The system displays detailed information as well as performance graphs for the volume.

Related reference

[Volume performance details](#) on page 156

You can view performance statistics of volumes from the Volume Performance page of the Reporting tab in the Element UI.

Volume performance details

You can view performance statistics of volumes from the Volume Performance page of the Reporting tab in the Element UI.

The following list describes the details that are available to you:

ID

The system-generated ID for the volume.

Name

The name given to the volume when it was created.

Account

The name of the account assigned to the volume.

Access Groups

The name of the volume access group or groups to which the volume belongs.

Volume Utilization

A percentage value that describes how much the client is using the volume.

Possible values:

- 0 = Client is not using the volume
- 100 = Client is using the max
- >100 = Client is using the burst

Total IOPS

The total number of IOPS (read and write) currently being executed against the volume.

Read IOPS

The total number of read IOPS currently being executed against the volume.

Write IOPS

The total number of write IOPS currently being executed against the volume.

Total Throughput

The total amount of throughput (read and write) currently being executed against the volume.

Read Throughput

The total amount of read throughput currently being executed against the volume.

Write Throughput

The total amount of write throughput currently being executed against the volume.

Total Latency

The average time, in microseconds, to complete read and write operations to a volume.

Read Latency

The average time, in microseconds, to complete read operations to the volume in the last 500 milliseconds.

Write Latency

The average time, in microseconds, to complete write operations to a volume in the last 500 milliseconds.

Queue Depth

The number of outstanding read and write operations to the volume.

Average IO Size

Average size in bytes of recent I/O to the volume in the last 500 milliseconds.

Viewing iSCSI sessions

You can view the iSCSI sessions that are connected to the cluster. You can filter the information to include only the desired sessions.

Steps

1. In the Element UI, select **Reporting > iSCSI Sessions**.
2. To see the filter criteria fields, click **Filter**.

Related reference

[iSCSI session details](#) on page 157

You can view information about the iSCSI sessions that are connected to the cluster.

iSCSI session details

You can view information about the iSCSI sessions that are connected to the cluster.

The following list describes the information that you can find about the iSCSI sessions:

Node

The node hosting the primary metadata partition for the volume.

Account

The name of the account that owns the volume. If value is blank, a dash (-) is displayed.

Volume

The volume name identified on the node.

Volume ID

ID of the volume associated with the Target IQN.

Initiator ID

A system-generated ID for the initiator.

Initiator Alias

An optional name for the initiator that makes finding the initiator easier when in a long list.

Initiator IP

The IP address of the endpoint that initiates the session.

Initiator IQN

The IQN of the endpoint that initiates the session.

Target IP

The IP address of the node hosting the volume.

Target IQN

The IQN of the volume.

Created On

Date the session was established.

Viewing Fibre Channel sessions

You can view the Fibre Channel (FC) sessions that are connected to the cluster. You can filter information to include only those connections you want displayed in the window.

Steps

1. In the Element UI, select **Reporting > FC Sessions**.
2. To see the filter criteria fields, click **Filter**.

Related reference

[Fibre Channel session details](#) on page 158

You can find information about the active Fibre Channel (FC) sessions that are connected to the cluster.

Fibre Channel session details

You can find information about the active Fibre Channel (FC) sessions that are connected to the cluster.

The following list describes the information you can find about the FC sessions connected to the cluster:

Node ID

The node hosting the session for the connection.

Node Name

System-generated node name.

Initiator ID

A system-generated ID for the initiator.

Initiator WWPN

The initiating worldwide port name.

Initiator Alias

An optional name for the initiator that makes finding the initiator easier when in a long list.

Target WWPN

The target worldwide port name.

Volume Access Group

Name of the volume access group that the session belongs to.

Volume Access Group ID

System-generated ID for the access group.

Troubleshooting drives

You can replace a failed solid-state drive (SSD) with a replacement drive. SSDs for SolidFire storage nodes are hot-swappable. If you suspect an SSD has failed, contact NetApp Support to verify the failure and walk you through the proper resolution procedure. NetApp Support also works with you to get a replacement drive according to your service-level agreement.

Hot-swappable in this case means that you can remove a failed drive from an active node and replace it with a new SSD drive from NetApp. It is not recommended that you should remove non-failed drives on an active cluster.

You should maintain on-site spares suggested by NetApp Support to allow for immediate replacement of the drive if it fails.

Note: For testing purposes, if you are simulating a drive failure by pulling a drive from a node, you must wait 30 seconds before inserting the drive back into the drive slot.

If a drive fails, Double Helix redistributes the data on the drive across the nodes remaining on the cluster. Multiple drive failures on the same node are not an issue since Element software protects against two copies of data residing on the same node. A failed drive results in the following events:

- Data is migrated off of the drive.
- Overall cluster capacity is reduced by the capacity of the drive.
- Double Helix data protection ensures that there are two valid copies of the data.



Attention: SolidFire storage systems do not support removal of a drive if it results in an insufficient amount of storage to migrate data.

Related concepts

[Basic MDSS drive troubleshooting](#) on page 160

You can recover metadata (or slice) drives by adding them back to the cluster in the event that one or both metadata drives fail. You can perform the recovery operation in the NetApp Element UI if the MDSS feature is already enabled on the node.

Related tasks

[Removing failed drives from the cluster](#) on page 160

The SolidFire system puts a drive in a failed state if the drive's self-diagnostics tells the node it has failed or if communication with the drive stops for five and a half minutes or longer. The system displays a list of the failed drives. You must remove a failed drive from the failed drive list in NetApp Element software.

[Removing MDSS drives](#) on page 162

You can remove the multi-drive slice service (MDSS) drives. This procedure applies only if the node has multiple slice drives.

Related information

[Replacing drives for SolidFire storage nodes](#)
[Replacing drives for H600S series storage nodes](#)

Removing failed drives from the cluster

The SolidFire system puts a drive in a failed state if the drive's self-diagnostics tells the node it has failed or if communication with the drive stops for five and a half minutes or longer. The system displays a list of the failed drives. You must remove a failed drive from the failed drive list in NetApp Element software.

About this task

Drives in the **Alerts** list show as **blockServiceUnhealthy** when a node is offline. When restarting the node, if the node and its drives come back online within five and a half minutes, the drives automatically update and continue as active drives in the cluster.

Steps

1. In the Element UI, select **Cluster > Drives**.
2. Click **Failed** to view the list of failed drives.
3. Note the slot number of the failed drive.
You need this information to locate the failed drive in the chassis.
4. Remove the failed drives using one of the following methods:

Option	Steps
To remove individual drives	<ol style="list-style-type: none">a. Click Actions for the drive you want to remove.b. Click Remove.
To remove multiple drives	<ol style="list-style-type: none">a. Select all the drives you want to remove, and click Bulk Actions.b. Click Remove.

Basic MDSS drive troubleshooting

You can recover metadata (or slice) drives by adding them back to the cluster in the event that one or both metadata drives fail. You can perform the recovery operation in the NetApp Element UI if the MDSS feature is already enabled on the node.

If either or both of the metadata drives in a node experiences a failure, the slice service will shut down and data from both drives will be backed up to different drives in the node.

The following scenarios outline possible failure scenarios, and provide basic recommendations to correct the issue:

System slice drive fails

- In this scenario, the slot 2 is verified and returned to an available state.
- The system slice drive must be repopulated before the slice service can be brought back online.
- You should replace the system slice drive, when the system slice drive becomes available, add the drive and the slot 2 drive at the same time.

Note: You cannot add the drive in slot 2 by itself as a metadata drive. You must add both drives back to the node at the same time.

Slot 2 fails

- In this scenario, the system slice drive is verified and returned to an available state.
- You should replace slot 2 with a spare, when slot 2 becomes available, add the system slice drive and the slot 2 drive at the same time.

System slice drive and slot 2 fails

- You should replace both system slice drive and slot 2 with a spare drive. When both drives become available, add the system slice drive and the slot 2 drive at the same time.

Order of operations

- Replace the failed hardware drive with a spare drive (replace both drives if both have failed).
- Add drives back to the cluster when they have been repopulated and are in an available state.

Verify operations

- Verify that the drives in slot 0 (or internal) and slot 2 are identified as metadata drives in the Active Drives list.
- Verify that all slice balancing has completed (there are no further moving slices messages in the event log for at least 30 minutes).

Related tasks

[Adding MDSS drives](#) on page 161

You can add a second metadata drive on a SolidFire node by converting the block drive in slot 2 to a slice drive. This is accomplished by enabling the multi-drive slice service (MDSS) feature. To enable this feature, you must contact NetApp Support.

Adding MDSS drives

You can add a second metadata drive on a SolidFire node by converting the block drive in slot 2 to a slice drive. This is accomplished by enabling the multi-drive slice service (MDSS) feature. To enable this feature, you must contact NetApp Support.

About this task

Getting a slice drive into an available state might require replacing a failed drive with a new or spare drive. You must add the system slice drive at the same time you add the drive for slot 2. If you try to add the slot 2 slice drive alone or before you add the system slice drive, the system will generate an error.

Steps

1. Click **Cluster > Drives**.
2. Click **Available** to view the list of available drives.
3. Select the slice drives to add.
4. Click **Bulk Actions**.
5. Click **Add**.
6. Confirm from the **Active Drives** tab that the drives have been added.

Removing MDSS drives

You can remove the multi-drive slice service (MDSS) drives. This procedure applies only if the node has multiple slice drives.

About this task

Note: If the system slice drive and the slot 2 drive fail, the system will shutdown slice services and remove the drives. If there is no failure and you remove the drives, both drives must be removed at the same time.

Steps

1. Click **Cluster > Drives**.
2. From the **Available** drives tab, click the check box for the slice drives being removed.
3. Click **Bulk Actions**.
4. Click **Remove**.
5. Confirm the action.

Troubleshooting nodes

You can remove nodes from a cluster for maintenance or replacement. You should use the NetApp Element UI or API to remove nodes before taking them offline.

An overview of the procedure to remove storage nodes is as follows:

- Ensure that there is sufficient capacity in the cluster to create a copy of the data on the node.
- Remove drives from the cluster by using the UI or the `RemoveDrives` API method. This results in the system migrating data from the node's drives to other drives in the cluster. The time this process takes is dependent on how much data must be migrated.
- Remove the node from the cluster.

Keep the following considerations in mind before you power down or power up a node:

- Powering down nodes and clusters involves risks if not performed properly. Powering down a node should be done under the direction of NetApp Support.
- If a node has been down longer than 5.5 minutes under any type of shutdown condition, Double Helix data protection begins the task of writing single replicated blocks to another node to replicate the data. In this case, contact NetApp Support for help with analyzing the failed node.
- To safely reboot or power down a node, you can use the `Shutdown` API command.
- If a node is in a down, or in an off state, you must contact NetApp Support before bringing it back online.
- After a node is brought back online, you must add the drives back to the cluster, depending on how long it has been out of service.

Related information

[Replacing a failed SolidFire chassis](#)

[Replacing a failed H600S series node](#)

Powering down a cluster

You can power down an entire cluster after you have contacted NetApp Support and completed preliminary steps.

Before you begin

Prepare the cluster for shutdown by doing the following:

- Stop all I/O.
- Disconnect all iSCSI sessions.

Steps

1. Navigate to the management virtual IP (MVIP) address on the cluster to open the Element UI.
2. Note the nodes listed in the Nodes list.
3. Run the `Shutdown` API method with the `halt` option specified on each Node ID in the cluster.

Working with per-node utilities for storage nodes

You can use the per-node utilities to troubleshoot network problems if the standard monitoring tools in the NetApp Element software UI do not give you enough information for troubleshooting. Per-node utilities provide specific information and tools that can help you troubleshoot network problems between nodes or with the management node.

Related tasks

[Accessing per-node settings using the per-node UI](#) on page 163

You can access network settings, cluster settings, and system tests and utilities in the per-node user interface after you enter the management node IP and authenticate.

[Running system tests using the per-node UI](#) on page 166

You can test changes to the network settings after you commit them to the network configuration. You can run the tests to ensure that the storage node is stable and can be brought online without any issues.

[Running system utilities using the per-node UI](#) on page 167

You can use the per-node UI for the storage node to create or delete support bundles, reset configuration settings for drives, and restart network or cluster services.

Related reference

[Network settings details from the per-node UI](#) on page 164

You can change the storage node network settings to give the node a new set of network attributes.

[Cluster settings details from the per-node UI](#) on page 165

You can verify cluster settings for a storage node after cluster configuration and modify the node hostname.

Accessing per-node settings using the per-node UI

You can access network settings, cluster settings, and system tests and utilities in the per-node user interface after you enter the management node IP and authenticate.

About this task

If you want to modify settings of a node in an `Active` state that is part of a cluster, you must log in as a cluster administrator user.

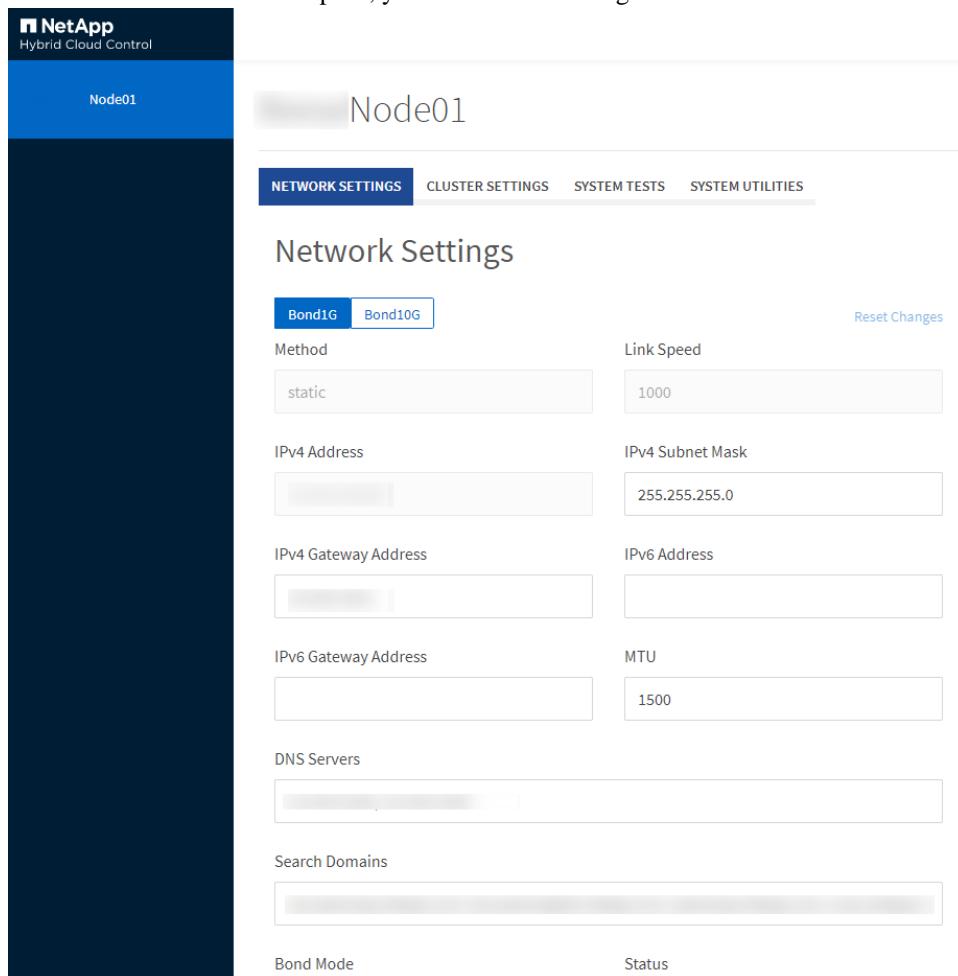
Tip: You should configure or modify one node at a time. You should ensure that the network settings specified are having the expected effect, and that the network is stable and performing well before you make modifications to another node.

Step

Open the per-node UI using one of the following methods:

- Enter the management IP address followed by `:442` in a browser window, and log in using an admin user name and password.
- In the Element UI, select **Cluster > Nodes**, and click the management IP address link for the node you want to configure or modify.

In the browser window that opens, you can edit the settings of the node.



The screenshot shows the Network Settings page for Node01. The top navigation bar includes the NetApp Hybrid Cloud Control logo, the node name 'Node01', and tabs for NETWORK SETTINGS, CLUSTER SETTINGS, SYSTEM TESTS, and SYSTEM UTILITIES. The NETWORK SETTINGS tab is selected. Below the tabs, the title 'Network Settings' is displayed. Under the 'Bond Mode' section, 'Bond1G' is selected. The 'Method' field is set to 'static'. The 'Link Speed' is set to '1000'. The 'IPv4 Address' and 'IPv4 Subnet Mask' fields are populated. The 'IPv4 Gateway Address' and 'IPv6 Address' fields are empty. The 'IPv6 Gateway Address' and 'MTU' fields are empty. The 'MTU' value is '1500'. The 'DNS Servers' and 'Search Domains' fields are empty. At the bottom, the 'Bond Mode' is 'Bond1G' and the 'Status' is 'Pending'.

Network settings details from the per-node UI

You can change the storage node network settings to give the node a new set of network attributes.

You can see the network settings for a storage node on the **Network Settings** page when you log in to the node (<https://<node IP>:442/hcc/node/network-settings>). You can select either **Bond1G** (management) or **Bond10G** (storage) settings. The following list describes the settings that you can modify when a storage node is in Available, Pending, or Active state:

Method

The method used to configure the interface. Possible methods:

- **loopback**: Used to define the IPv4 loopback interface.
- **manual**: Used to define interfaces for which no configuration is done by default.
- **dhcp**: Used to obtain an IP address via DHCP.
- **static**: Used to define Ethernet interfaces with statically allocated IPv4 addresses.

Link Speed

The speed negotiated by the virtual NIC.

IPv4 Address

The IPv4 address for the eth0 network.

IPv4 Subnet Mask

Address subdivisions of the IPv4 network.

IPv4 Gateway Address

Router network address to send packets out of the local network.

IPv6 Address

The IPv6 address for the eth0 network.

IPv6 Gateway Address

Router network address to send packets out of the local network.

MTU

Largest packet size that a network protocol can transmit. Must be greater than or equal to 1500. If you add a second storage NIC, the value should be 9000.

DNS Servers

Network interface used for cluster communication.

Search Domains

Search for additional MAC addresses available to the system.

Bond Mode

Can be one of the following modes:

- ActivePassive (default)
- ALB
- LACP

Status

Possible values:

- UpAndRunning
- Down
- Up

Virtual Network Tag

Tag assigned when the virtual network was created.

Routes

Static routes to specific hosts or networks via the associated interface the routes are configured to use.

Cluster settings details from the per-node UI

You can verify cluster settings for a storage node after cluster configuration and modify the node hostname.

The following list describes the cluster settings for a storage node indicated from the **Cluster Settings** page of the per-node UI (<https://<node IP>:442/hcc/node/cluster-settings>).

Role

Role the node has in the cluster. Possible values:

- Storage: Storage or Fibre Channel node.
- Management: Node is a management node.

Hostname

Name of the node.

Cluster

Name of the cluster.

Cluster Membership

State of the node. Possible values:

- **Available:** The node has no associated cluster name and is not yet part of a cluster.
- **Pending:** The node is configured and can be added to a designated cluster. Authentication is not required to access the node.
- **PendingActive:** The system is in the process of installing compatible software on the node. When complete, the node will move to the Active state.
- **Active:** The node is participating in a cluster. Authentication is required to modify the node.

Version

Version of the Element software running on the node.

Ensemble

Nodes that are part of the database ensemble.

Node ID

ID assigned when a node is added to the cluster.

Cluster Interface

Network interface used for cluster communication.

Management Interface

Management network interface. This defaults to Bond1G but can also use Bond10G.

Storage Interface

Storage network interface using Bond10G.

Encryption Capable

Indicates whether or not the node supports drive encryption.

Running system tests using the per-node UI

You can test changes to the network settings after you commit them to the network configuration. You can run the tests to ensure that the storage node is stable and can be brought online without any issues.

Before you begin

You have logged in to the per-node UI for the storage node.

Steps

1. Click **System Tests**.
2. Click **Run Test** next to the test you want to run or select **Run All Tests**.

Note: Running all test operations can be time consuming and should be done only at the direction of NetApp Support.

Test Connected Ensemble

Tests and verifies the connectivity to a database ensemble. By default, the test uses the ensemble for the cluster the node is associated with. Alternatively you can provide a different ensemble to test connectivity.

Test Connect Mvip

Pings the specified management virtual IP (MVIP) address and then executes a simple API call to the MVIP to verify connectivity. By default, the test uses the MVIP for the cluster the node is associated with.

Test Connect Svip

Pings the specified storage virtual IP (SVIP) address using Internet Control Message Protocol (ICMP) packets that match the Maximum Transmission Unit (MTU) size set on the network adapter. It then connects to the SVIP as an iSCSI initiator. By default, the test uses the SVIP for the cluster the node is associated with.

Test Hardware Config

Tests that all hardware configurations are correct, validates firmware versions are correct, and confirms all drives are installed and running properly. This is the same as factory testing.

Note: This test is resource intensive and should only be run if requested by NetApp Support.

Test Local Connectivity

Tests the connectivity to all of the other nodes in the cluster by pinging the cluster IP (CIP) on each node. This test will only be displayed on a node if the node is part of an active cluster.

Test Locate Cluster

Validates that the node can locate the cluster specified in the cluster configuration.

Test Network Config

Verifies that the configured network settings match the network settings being used on the system. This test is not intended to detect hardware failures when a node is actively participating in a cluster.

Test Ping

Pings a specified list of hosts or, if none are specified, dynamically builds a list of all registered nodes in the cluster and pings each for simple connectivity.

Test Remote Connectivity

Tests the connectivity to all nodes in remotely paired clusters by pinging the cluster IP (CIP) on each node. This test will only be displayed on a node if the node is part of an active cluster.

Running system utilities using the per-node UI

You can use the per-node UI for the storage node to create or delete support bundles, reset configuration settings for drives, and restart network or cluster services.

Before you begin

You have logged in to the per-node UI for the storage node.

Steps

1. Click **System Utilities**.
2. Click the button for the system utility that you want to run.

Control Power

Reboots, power cycles, or shuts down the node.



Attention: This operation causes temporary loss of networking connectivity.

Specify the following parameters:

- Action: Options include `Restart` and `Halt` (power off).
- Wakeup Delay: Any additional time before the node comes back online.

Collect Node Logs

Creates a support bundle under the node's `/tmp/bundles` directory.

Specify the following parameters:

- Bundle Name: Unique name for each support bundle created. If no name is provided, then "supportbundle" and the node name are used as the file name.
- Extra Args: This parameter is fed to the `sf_make_support_bundle` script. This parameter should be used only at the request of NetApp Support.
- Timeout Sec: Specify the number of seconds to wait for each individual ping response.

Delete Node Logs

Deletes any current support bundles on the node that were created using **Create Cluster Support Bundle** or the `CreateSupportBundle` API method.

Reset Drives

Initializes drives and removes all data currently residing on the drive. You can reuse the drive in an existing node or in an upgraded node.

Specify the following parameter:

- Drives: List of device names (not driveIDs) to reset.

Reset Network Config

Helps resolve network configuration issues for an individual node and resets an individual node's network configuration to the factory default settings.

Reset Node

Resets a node to the factory settings. All data is removed but network settings for the node are preserved during this operation. Nodes can only be reset if they are unassigned to a cluster and in Available state.



Attention: All data, packages (software upgrades), configurations, and log files are deleted from the node when you use this option.

Restart Networking

Restarts all networking services on a node.



Attention: This operation can cause temporary loss of network connectivity.

Restart Services

Restarts Element software services on a node.



Attention: This operation can cause temporary node service interruption. You

should perform this operation only at the direction of NetApp Support.

Specify the following parameters:

- Service: Service name to be restarted.
- Action: Action to perform on the service. Options include `start`, `stop` and `restart`.

Working with the management node

You can use the management node (mNode) to upgrade system services, manage cluster assets and settings, run system tests and utilities, configure Active IQ for system monitoring, and enable NetApp Support access for troubleshooting.

For management node information, see management node documentation.

[Management node overview](#)

Understanding cluster fullness levels

The cluster running Element software generates cluster faults to warn the storage administrator when the cluster is running out of capacity. There are three levels of cluster fullness, all of which are displayed in the NetApp Element UI: warning, error, and critical.

The system uses the `BlockClusterFull` error code to warn about cluster block storage fullness. You can view the cluster fullness severity levels from the Alerts tab of the Element UI.

The following list includes information about the `BlockClusterFull` severity levels:

Warning

This is a customer-configurable warning that appears when the cluster's block capacity is approaching the error severity level. By default, this level is set at three percent under the error level and can be tuned via the Element UI and API. You must add more capacity, or free up capacity as soon as possible.

Error

When the cluster is in this state, if a node is lost, there will not be enough capacity in the cluster to rebuild Double Helix data protection. New volume creation, clones, and snapshots are all blocked while the cluster is in this state. This is not a safe or recommended state for any cluster to be in. You must add more capacity, or free up capacity immediately.

Critical

This critical error has occurred because the cluster is 100 percent consumed. It is in a read-only state and no new iSCSI connections can be made to the cluster. When this stage is reached, you must free up or add more capacity immediately.

The system uses the `MetadataClusterFull` error code to warn about cluster metadata storage fullness. You can view the cluster metadata storage fullness from the Cluster Capacity section on the Overview page of the Reporting tab in the Element UI.

The following list includes information about the `MetadataClusterFull` severity levels:

Warning

This is a customer-configurable warning that appears when the cluster's metadata capacity is approaching the error severity level. By default, this level is set at three percent under the error level and can be tuned via the Element API. You must add more capacity, or free up capacity as soon as possible.

Error

When the cluster is in this state, if a node is lost, there will not be enough capacity in the cluster to rebuild Double Helix data protection. New volume creation, clones, and snapshots are all blocked while the cluster is in this state. This is not a safe or recommended state for any cluster to be in. You must add more capacity, or free up capacity immediately.

Critical

This critical error has occurred because the cluster is 100 percent consumed. It is in a read-only state and no new iSCSI connections can be made to the cluster. When this stage is reached, you must free up or add more capacity immediately.

Note: The following applies to two-node cluster thresholds:

- Metadata fullness error is 20% below critical.
- Block fullness error is 1 block drive (including stranded capacity) below critical; meaning that it is two block drives worth of capacity below critical.

Contacting NetApp Support

If you need help with or have questions or comments about NetApp products, contact NetApp Support.

- Web:
mysupport.netapp.com
- Phone:
 - 888.4.NETAPP (888.463.8277) (US and Canada)
 - 00.800.44.638277 (EMEA/Europe)
 - +800.800.80.800 (Asia/Pacific)

Where to find product documentation and other information

You can learn more about using and managing NetApp HCI and SolidFire all-flash storage from the resources available in the Documentation Centers and Resources pages for both products.

In the Documentation Centers, you can also find information about hardware installation and maintenance, additional content resources available, links to known issues and resolved issues, and the latest release notes. On the Resources pages, you can find links to data sheets, technical reports, white papers, and videos.

- [*NetApp HCI Documentation*](#)
- [*NetApp HCI Documentation Center*](#)
- [*NetApp HCI Resources page*](#)
- [*SolidFire and Element 12.0 Documentation Center*](#)
- [*SolidFire and Element 11.8 Documentation Center*](#)
- [*SolidFire and Element 11.7 Documentation Center*](#)
- [*SolidFire and Element 11.5 Documentation Center*](#)
- [*SolidFire and Element 11.3 Documentation Center*](#)
- [*SolidFire Resources page*](#)

Copyright and trademark

Copyright

Copyright © 2020 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>