



NetApp SolidFire Element OS

User Guide

Version 10.3

June 2018 | 215-13201_A0
doccomments@netapp.com

 **NetApp**[®]

Contents

About this guide	8
SolidFire system overview	9
Clusters	9
Nodes	9
Storage nodes	9
Fibre Channel nodes	10
Networking	10
Switch configuration for SolidFire clusters	10
Network port requirements	11
System setup	14
Installing a management node	14
Setting up the connection.json file for Active IQ	15
Help options for the manage-collector.py utility	16
Configuring a storage node	17
Configuring the node using the node UI	18
Configuring the node using the TUI	19
Configuring iDRAC on each new node	19
Configuring a Fibre Channel node	19
Creating a new cluster	20
Adding drives to a cluster	21
Changing the Element OS default SSL certificate	22
Using the Element OS web UI	23
Using filters	23
Sorting lists	24
Using the API log	24
Web UI and cluster load	25
Icons in the Element OS web UI	25
Providing feedback	26
System management	27
Managing cluster administrator users	27
SolidFire storage user types	27
Cluster Admins details	28
Creating a cluster admin account	28
Editing cluster admin permissions	29
Deleting a cluster admin account	30
Changing passwords for cluster administrator accounts	30
Managing LDAP	30
Configuring cluster settings	32
Setting cluster full threshold	32
Encryption at rest	33
Enabling and disabling encryption for a cluster	33

Managing Terms of Use	34
Network time protocol	35
Enabling a broadcast client	35
Managing SNMP	35
Managing drives	37
Managing nodes	38
Viewing Fibre Channel ports details	41
Fibre Channel ports details	41
Managing virtual networks	42
Upgrading Element OS software	45
Using HealthTools for software upgrades (Element OS 10.2 and later)	45
Upgrading management node software	49
Upgrading Element OS 10.2 and later	50
Upgrading Element OS 10.2 and later on dark sites	51
Upgrading software for versions earlier than Element OS 10.2	52
Data management	54
Working with accounts	54
Creating an account	54
Account details	55
Viewing individual account details	55
Editing an account	55
Deleting an account	56
Working with volumes	56
SolidFire Quality of Service	57
QoS policies	59
Creating a volume	61
Volume details	61
Viewing individual volume performance details	62
Editing active volumes	63
Deleting a volume	64
Restoring a deleted volume	65
Purging a volume	65
Cloning a volume	65
Assigning LUNs to Fibre Channel volumes	67
Applying a QoS policy to volumes	67
Removing the QoS policy association of a volume	67
Working with virtual volumes	68
Enabling virtual volumes	68
Viewing virtual volume details	69
Virtual volume details	70
Individual virtual volume details	71
Deleting a virtual volume	72
Storage containers	73
Protocol endpoints	75
Bindings	76

Bindings details	76
Host details	76
Working with access groups and initiators	77
Creating an access group	77
Volume access group details	78
Viewing individual access group details	79
Adding volumes to an access group	79
Removing volumes from an access group	79
Creating an initiator	80
Editing an initiator	81
Adding a single initiator to an access group	81
Adding multiple initiators to an access group	82
Removing initiators from an access group	82
Deleting an access group	83
Deleting an initiator	83
Data protection	84
Using individual volume snapshots	84
Creating a volume snapshot	84
Editing snapshot retention	85
Deleting a snapshot	86
Cloning a volume from a snapshot	86
Rolling back a volume to a snapshot	87
Volume snapshot backup operations	87
Using group snapshots	89
Group snapshot details	89
Creating a group snapshot	90
Editing group snapshots	91
Deleting a group snapshot	91
Rolling back volumes to a group snapshot	92
Editing members of group snapshot	92
Cloning multiple volumes	93
Cloning multiple volumes from a group snapshot	93
Using snapshot schedules	94
Creating a snapshot schedule	94
Editing a snapshot schedule	95
Deleting a snapshot schedule	96
Copying a snapshot schedule	96
Using replication between SolidFire clusters	96
Configuring cluster and volume pairing for real-time replication	96
Cluster pairs	97
Pairing clusters	97
Cluster pair details	99
Deleting a cluster pair	100
Volume pairs	100
Volume pair details	104

Editing volume pairs	107
Deleting volume pairs	107
Using replication between SolidFire and ONTAP clusters	107
SnapMirror overview	108
Enabling SnapMirror on the cluster	108
Enabling SnapMirror on the volume	109
SnapMirror endpoints	109
SnapMirror labels	111
SnapMirror relationships	112
Disaster recovery using SnapMirror	116
Backing up and restoring volumes	120
Backing up a volume to an Amazon S3 object store	120
Backing up a volume to an OpenStack Swift object store	121
Backing up a volume to a SolidFire cluster	121
Restoring a volume from backup on an Amazon S3 object store	122
Restoring a volume from backup on an OpenStack Swift object store	123
Restoring a volume from backup on a SolidFire cluster	124
System monitoring and troubleshooting	125
Viewing information about system events	125
Event types	126
Viewing status of running tasks	128
Viewing system alerts	128
Alert error codes	129
Viewing node performance activity	130
Viewing volume performance	131
Volume performance details	131
Viewing iSCSI sessions	132
iSCSI session details	133
Viewing Fibre Channel sessions	133
Fibre Channel session details	134
Troubleshooting drives	134
Removing failed drives	135
Basic MDSS drive troubleshooting	135
Adding MDSS drives	136
Removing MDSS drives	137
Replacing failed drives in the chassis	137
Troubleshooting nodes	137
Powering down a cluster	138
Working with per-node utilities	138
Per-node network settings details	139
Per-node cluster settings details	139
Running system tests	140
Running system utilities from the per-node UI	141
Working with the management node	142
Accessing a management node	142

Management node network settings	142
Management node cluster settings	143
Testing the management node settings	143
Running system utilities from the management node	144
Configuring a proxy server for the management node	144
Enabling remote NetApp Support connections	145
Understanding cluster fullness levels	146
Accessing per-node settings	147
Enabling FIPS 140-2 on your cluster	147
SSL ciphers	148
Enabling and disabling support access	149
Where to find additional information	150
Copyright information	151
Trademark information	152
How to send comments about documentation and receive update notifications	153
Index	154

About this guide

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

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 system overview

The SolidFire storage system contains hardware and software designed for complete system automation and management. A cluster serves as the storage hub and comprises four or more nodes. Double Helix data protection ensures that there are two copies of data on a cluster.

Using the SolidFire Element OS web UI, you can set up and monitor SolidFire cluster storage capacity and performance, and manage storage activity across a multi-tenant infrastructure. The Element OS web UI is built on the SolidFire API which enables you to see system adjustments almost immediately.

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 OS web 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.

Nodes

SF-series nodes are the hardware that are grouped into a cluster to be accessed as block storage. There are two types of SF-series nodes: storage nodes and Fibre Channel nodes.

Related concepts

[Storage nodes](#) on page 9

[Fibre Channel nodes](#) on page 10

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 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.

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 SolidFire clusters](#) on page 10

Related references

[Network port requirements](#) on page 11

Switch configuration for SolidFire clusters

The SolidFire system has certain switch requirements and best practices for optimal storage performance.

SolidFire storage clusters require 10GbE Ethernet switches for iSCSI storage services and node intra-cluster services communication. 1GbE switches can be used for the following 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

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

- Deploy a pair of 1GbE switches to provide high availability and load sharing for non-storage traffic in the cluster.

- Configure and utilize jumbo frames (an MTU size of 9216 bytes) on the storage network switches. 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.

All ports are TCP unless stated otherwise, and should permit bi-directional communications between the NetApp SolidFire Support Server, management node, and the SolidFire nodes.

Tip:

NetApp recommends enabling ICMP between the management node, SolidFire nodes, and cluster MVIP.

The following abbreviations are used in the table:

- MIP: Management IP address
- SIP: Storage IP address
- MVIP: Management virtual IP address
- SVIP: Storage virtual IP address

Source	Destination	Port	Description
iSCSI clients	Storage cluster MVIP	443	(Optional) UI and API access
iSCSI clients	Storage cluster SVIP	3260	Client iSCSI communications
iSCSI clients	Storage node SIP	3260	Client iSCSI communications
Management node	sfsupport.solidfire.com	22	Reverse SSH tunnel for support access
Management node	Storage node MIP	22	SSH access for support
Management node	DNS servers	53 TCP/UDP	DNS lookup
Management node	Storage node MIP	442	UI and API access to storage node
Management node	monitoring.solidfire.com	443	Storage cluster reporting to Active IQ
Management node	Storage cluster MVIP	443	UI and API access to storage cluster
SNMP server	Storage cluster MVIP	161 UDP	SNMP polling
SNMP server	Storage node MIP	161 UDP	SNMP polling
Storage node MIP	DNS servers	53 TCP/UDP	DNS lookup
Storage node MIP	Management node	80	Element OS updates
Storage node MIP	S3/Swift endpoint	80	(Optional) HTTP communication to S3/Swift endpoint for backup and recovery

Source	Destination	Port	Description
Storage node MIP	NTP server	123 TCP/UDP	NTP
Storage node MIP	Management node	162 UDP	SNMP traps
Storage node MIP	SNMP server	162 UDP	SNMP traps
Storage node MIP	LDAP server	389 TCP/UDP	LDAP lookup
Storage node MIP	Remote storage cluster MVIP	443	Remote replication cluster pairing communication
Storage node MIP	Remote storage node MIP	443	Remote replication cluster pairing communication
Storage node MIP	S3/Swift endpoint	443	(Optional) HTTPS communication to S3/Swift endpoint for backup and recovery
Storage node MIP	Management node	514 TCP/UDP	Syslog forwarding
Storage node MIP	Syslog server	514 TCP/UDP	Syslog forwarding
Storage node MIP	LDAPS server	636 TCP/UDP	LDAPS lookup
Storage node MIP	Remote storage node MIP	2181	Remote replication intercluster communication
Storage node SIP	S3/Swift endpoint	80	(Optional) HTTP communication to S3/Swift endpoint for backup and recovery
Storage node SIP	S3/Swift endpoint	443	(Optional) HTTPS communication to S3/Swift endpoint for backup and recovery
Storage node SIP	Remote storage node SIP	2181	Remote replication intercluster communication
Storage node SIP	Storage node SIP	3260	Internode iSCSI
Storage node SIP	Remote storage node SIP	4000 through 4020	Remote replication node-to- node data transfer
System administrator PC	Management node	442	HTTPS UI access to management node
System administrator PC	Storage node MIP	442	HTTPS UI and API access to storage node
System administrator PC	Management node	443	HTTPS UI and API access to management node
System administrator PC	Storage cluster MVIP	443	HTTPS UI and API access to storage cluster

Source	Destination	Port	Description
System administrator PC	Storage node MIP	443	HTTPS storage cluster creation, post-deployment UI access to storage cluster
vCenter Server	Storage cluster MVIP	443	vCenter Plug-in API access
vCenter Server	Management node	8080/8443	(Optional) vCenter Plug-in QoSSIOC service. 8080 redirects to 8443.
vCenter Server	Storage cluster MVIP	8444	vCenter VASA provider access (VVols only)
vCenter Server	Management node	9443	vCenter Plug-in registration; the port can be closed after registration is complete.

Related information

[VMware documentation](#)

System setup

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

The SolidFire 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 SolidFire 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 system, you must follow a specific order of operations to ensure that your nodes and clusters are configured correctly.

Steps

1. [Installing a management node](#) on page 14
2. [Setting up the connection.json file for Active IQ](#) on page 15
3. [Configuring a storage node](#) on page 17
4. [Configuring a Fibre Channel node](#) on page 19
5. [Creating a new cluster](#) on page 20
6. [Adding drives to a cluster](#) on page 21
7. [Changing the Element OS default SSL certificate](#) on page 22

Related concepts

[Using the Element OS web UI](#) on page 23

[Working with accounts](#) on page 54

[Working with volumes](#) on page 56

Installing a management node

You can install the management node using the appropriate image for your configuration. A management node is a virtual machine-based node used to upgrade the system software, connect to Active IQ for system monitoring, and allow NetApp Support to access your nodes if you need help troubleshooting a problem.

Before you begin

You have identified the management node image type that is correct for your platform. See the following table for guidance:

Platform	Installation image type
Microsoft Hyper-V	.iso
KVM	.iso
VMware vSphere	.iso, .ova
Citrix XenServer	.iso
OpenStack	.qcow2

About this task

You must use a minimum screen resolution of 1024 x 768 at 16-bit color.

Steps

1. Create a new 64-bit virtual machine with the following configuration:
 - One virtual CPU
 - 4GB RAM
 - 100GB virtual disk
 - One virtual network interface with internet access
 - (Optional) One virtual network interface with management network access to the SolidFire cluster
2. Attach the `solidfire-fdva-xxxx-xxxx.iso` to the virtual machine, and boot to the `.iso` install image.

Note:

- Access the NetApp Support Site for the latest version of the management node `.iso` image.
 - Installing a management node removes all data from the virtual machine (VM).
3. Power on the management node after the installation completes.
 4. Create a management node admin user using the terminal user interface (TUI).

Tip: To enter text, press **Enter** on the keyboard to open edit mode. After you enter text, press **Enter** again to close the edit mode. To navigate between fields, use the arrow keys.
 5. Configure the management node network using the TUI.

Related concepts

[Working with the management node](#) on page 142

Related tasks

[Accessing a management node](#) on page 142

[Configuring the node using the TUI](#) on page 19

Setting up the `connection.json` file for Active IQ

You must configure the `connection.json` file when connecting to SolidFire Active IQ from behind a proxy server.

About this task

You might choose to leave out the proxy and certificates steps of this procedure if they are not required in your environment. The `proxyUsername` and `proxyPassword` are optional even if you specify a proxy server. If you specify a `proxyIP`, you must also specify a `proxyPort`.

Note: The certificates option might be required if the collector is installed on something other than a management node. By default, the certificates option looks for the `/etc/ssl/certs/ca-certificates.crt` file to get the set of trusted root certificates to validate the remote support server SSL certification. If that file does not exist, you can use the certificates file that is

maintained by the cURL project. The certificates file is located at: <http://curl.haxx.se/ca/cacert.pem>. Save the `ca.crt` file in a desired location, and point the certificates option to that file.

Steps

1. Open a terminal window and use Secure Shell (SSH) to connect to your management node.
2. Become root with the following command:

```
sudo su
```

3. Change to the following directory: `cd / solidfire/collector`
4. Change the permissions for the `collector.py` file to `775` with the following command:

```
sudo chmod 775 collector.py
```

5. View the help to see the options that you can use to configure the `connection.json` file.

```
sudo ./manage-collector.py --help
```

6. Set the user name and MVIP for the collection configuration in the `connection.json` file.
See the following sample input command:

```
./manage-collector.py --set-username <username> --set-mvip <mvip>
```

The script automatically saves the `connection.json` file.

7. Set the password using the `set-password` command.
See the following sample input command:

```
./manage-collector.py --set-password <password>
```

Note: When you enter a password using the `set-password` command, you are prompted to enter the password and reenter it to confirm the password.

8. Restart the collector service with the following command:

```
sudo restart sfcollector
```

9. Verify the connection is working with the following command:

```
tail -f /var/log/sf-collector.log
```

Related references

[Help options for the `manage-collector.py` utility](#) on page 16

Help options for the `manage-collector.py` utility

You can view the full list of help options that can be used to configure the `connection.json` file.

-h --help

Show help message and exit.

--config CONFIG

Collector configuration file to manage (default: ./connection.json).

--save-config

Save the configuration to the collector configuration file. This option is not necessary when calling any of the set commands; the configuration is saved automatically when using those commands.

--set-username USERNAME

Set the cluster user name in the collector configuration file.

--set-password

Set the cluster password in the collector configuration file.

Note: The new password will be captured at the prompt.

--set-mvip MVIP

Set the cluster management virtual IP (MVIP) address in the collector configuration file.

--set-remoteHost REMOTEHOST

Set the remote host in the collector configuration file.

--get-all

Get all the parameters from the collector configuration file.

--get-username

Get the cluster user name from the collector configuration file.

--get-password

Get the cluster password from the collector configuration file.

--get-mvip

Get the cluster MVIP from the collector configuration file.

--get-remoteHost

Get the remote host from the collector configuration file.

--debug

Enable debug messages.

Related tasks

[Setting up the connection.json file for Active IQ](#) on page 15

Configuring a storage node

You must configure individual nodes before you can add them to a cluster. When you install and cable a node in a rack unit and power it on, the terminal user interface (TUI) displays the fields necessary to configure the node. Ensure that you have the necessary configuration information for the node before proceeding.

Alternatively, you can configure these settings by accessing the node via the Element OS web UI using the Dynamic Host Configuration Protocol (DHCP) 1G management IP address displayed in the TUI. The DHCP address is located in the menu bar at the top of the TUI.

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 web UI, TUI, or API. During this initial configuration, you can add the 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 SolidFire API methods.

Related concepts

[Configuring iDRAC on each new node](#) on page 19

Related tasks

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

[Configuring the node using the TUI](#) on page 19

[Creating a new cluster](#) on page 20

Related references

[Node states](#) on page 40

Related information

[SolidFire system installation and configuration \(video\)](#)

Configuring the node using the node UI

You can configure nodes using the Node Configuration user interface.

About this task

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 Bond1G and Bond10G 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 Settings page of the Element OS web 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, 172.25.103.6:442.
The Network Settings tab opens with the Network Settings – Bond1G section.
2. Enter the 1G network settings.
3. Click **Save Changes**.
4. Click **Bond10G** to display the settings for the 10G network settings.
5. Enter the 10G network settings.
6. Click **Save Changes**.
7. Click **Cluster Settings**.
8. Enter the hostname for the 10G network.
9. Click **Save Changes**.

Configuring the node using the TUI

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

About this task

Attention: You should configure the Bond1G and Bond10G 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 Settings page of the Element OS web UI.

Steps

1. Attach a keyboard and monitor to the node and then power on the node.
The TUI appears on the `tty1` terminal with the Network Settings tab. If a DHCP server is running on the network with available IP addresses, the 1GbE address appears in the Address field.
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. Use the on-screen navigation to configure the 1G and 10G network settings for the node.
Tip: To enter text, press **Enter** on the keyboard to open edit mode. After you enter text, press **Enter** again to close the edit mode. To navigate between fields, use the arrow keys.
3. Enter `s` to save the settings, and then enter `y` to accept the changes.
4. Enter `c` to navigate to the **Cluster** tab.
5. Use the on-screen navigation to configure the cluster settings for the node.
All the nodes in a cluster must have identical cluster names. Cluster names are case-sensitive.
6. Enter `s` to save the settings, and then enter `y` to accept the changes.
The node is put in a pending state and can be added to an existing cluster or a new cluster.

Related information

[How to set up and install a SolidFire cluster in less than five minutes \(video\)](#)

Configuring iDRAC on each new node

NetApp installs Dell iDRAC Enterprise on each node. iDRAC enables you to remotely manage and monitor the underlying hardware of each node in the cluster.

Related information

[How to configure iDRAC on a SolidFire storage node](#)

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 OS 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 allow for the best overall performance from the Fibre Channel systems.

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 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 OS web UI, you can find the WWPN addresses from the WWPN column of the FC Ports tab, which you access from the Cluster tab.

Related concepts

[Configuring iDRAC on each new node](#) on page 19

Related tasks

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

[Configuring the node using the TUI](#) on page 19

[Creating a new cluster](#) on page 20

Related information

[SolidFire Fibre Channel Configuration Guide](#)

Creating a new cluster

You can create a new cluster after you have configured 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 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.

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**, if you approve.
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.

Adding drives to a cluster

You will be directed to the Available Drives list when drives have been detected when you create a cluster. You will be prompted to add all the available drives at that time.

About this task

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 using either the web UI or API before it can participate in the cluster.

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. Select **Cluster > Drives**.
2. Click **Available** to view the list of available drives.
3. Add drives by following one of the following methods:
 - To add individual drives, click the Actions icon for the drive you want to add.
 - To add multiple drives, select the check box of the drives to add, and then click Bulk Actions.
4. Click **Add**.

Changing the Element OS default SSL certificate

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

When an Element OS 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 OS web UI, per-node UI, or APIs. Element OS 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 details about each method, see the *NetApp SolidFire Element OS 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.

Related information

[SolidFire Element OS API Reference Guide](#)

Using the Element OS web UI

You can use the Element OS web user interface (UI) to monitor and perform common tasks on your SolidFire system. You can access the UI by using the management virtual IP (MVIP) address of the primary cluster node.

Before you access the UI, you must ensure that popup blockers and NoScript settings are disabled in your browser.

You can access the UI by entering the following URL in your browser window:

```
https://<MVIP formatted as 1-2-3-4>.ip.solidfire.net
```

If your browser cannot connect to the Internet, you can enter the MVIP address as a URL in your browser window:

```
http://<MVIP address>
```

You should click through any authentication certificate messages.

Related tasks

[Using filters](#) on page 23

[Using the API log](#) on page 24

[Sorting lists](#) on page 24

[Providing feedback](#) on page 26

Related references

[Icons in the Element OS web UI](#) on page 25

Using filters

You can sort and filter list information on pages in the Element OS web UI. When viewing lists (such as volumes, snapshots, and so on), you can use filter functionality to focus the information and make it more easily fit on the screen.

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, perform the following steps:
 - a. Click **Add**.
 - b. Follow Steps 3 through 6 to add another filter.

- 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 OS web UI. This helps you arrange the information you need on the screen.

Steps

- To sort on a single column, click the column heading until the information is sorted in the desired order.
- 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.

- To reorder the sort criteria, perform the following steps:
 - Click **Sort**.
The system populates the Sort By field with your column selections.
 - Arrange the columns in the Sort By field in the order you want the list to be sorted.
The system sorts the list information.
- To remove a single sort criterion, click the **Remove** icon next to the name of a sort criteria.
- Optional: To remove all sort criteria, click **Clear All**.

Using the API log

The SolidFire system uses the SolidFire API as the foundation for its features and functionality. The Element OS web 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.

Steps

- From the SolidFire Element OS web UI navigation bar, click **API Log**.
- To modify the type of API activity displayed in the **API Log** window, perform the following steps:
 - Select **Requests** to display API request traffic.
 - Select **Responses** to display API response traffic.
 - 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

[SolidFire Element OS API Reference Guide](#)

Web UI and cluster load

Depending on API response times, the cluster might automatically adjust the data refresh interval for certain portions of the 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 web 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 web 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 OS web UI

You can find information about the icons in the Element OS web UI.

The following table provides a quick reference:

Icon	Description
	Actions
	Backup to
	Clone or copy
	Delete or purge
	Edit
	Filter
	Pair
	Refresh
	Restore
	Restore from
	Rollback
	Snapshot

Providing feedback

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

Steps

1. From any page in the 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 SolidFire system in the Element OS web UI. This includes creating and managing cluster administrators, managing cluster settings, and upgrading software.

Related concepts

[Managing cluster administrator users](#) on page 27

[Configuring cluster settings](#) on page 32

[Upgrading Element OS software](#) on page 45

Managing cluster administrator users

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

Related concepts

[SolidFire storage user types](#) on page 27

Related tasks

[Creating a cluster admin account](#) on page 28

[Editing cluster admin permissions](#) on page 29

[Deleting a cluster admin account](#) on page 30

[Changing passwords for cluster administrator accounts](#) on page 30

[Configuring LDAP](#) on page 31

[Disabling LDAP](#) on page 32

Related references

[Cluster Admins details](#) on page 28

SolidFire storage user types

There are two types of administrators that can exist in a SolidFire cluster: the primary cluster administrator account and a cluster admin role.

The following types of administrators can exist in a cluster:

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. Two (or more) cluster administrators with administrator access permissions must exist before you can delete the primary cluster administrator account. You can change the password for this administrator account.

Cluster admin account

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

Note: A local (non-LDAP) cluster admin 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 admin account.

Username

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

Access

Shows 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

Shows the type of cluster administrator. Possible values:

- Cluster
- Ldap

Attributes

If the cluster admin user was created using the SolidFire API, this column shows any name-value pairs that were set using that method.

Related information

[SolidFire Element OS API Reference Guide](#)

Creating a cluster admin 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.

Steps

1. Click **Users > Cluster Admins**.
2. Click **Create Cluster Admin**.
3. To create a cluster-wide (non-LDAP) cluster administrator account, perform the following actions:
 - a. Select the **Cluster** user type.
 - b. Enter a user name.
 - c. Enter a password for the account.
 - d. Confirm the 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**.
4. To create a cluster administrator account in the LDAP directory, perform the following actions:
 - a. Select the **LDAP** user type.
 - b. Follow the example in the text box to enter a full distinguished name for the user.
 - c. Select user permissions to apply to the account.
 - d. Select the check box to agree to the End User License Agreement.
 - e. Click **Create Cluster Admin**.

Related tasks

[Configuring LDAP](#) on page 31

Editing cluster admin 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 admin you want to edit.
3. Click **Edit**.
4. Select user permissions to apply to the account.
5. Click **Save Changes**.

Deleting a cluster admin account

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

Steps

1. Click **Users > Cluster Admins**.
2. Click the Actions icon for the cluster admin you want to delete.
3. In the resulting menu, select **Delete**.
4. Confirm the action.

Changing passwords for cluster administrator accounts

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

Steps

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

Related concepts

[SolidFire storage user types](#) on page 27

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.

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.

LDAP Authentication Enabled

To configure LDAP, you must enable authentication.

LDAP Servers

Address of an LDAP or LDAPS directory server.

Auth Type

Identifies which user authentication method is used. Possible values:

- DirectBind
- SearchAndBind

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.

Group Search Type

Controls the default group search filter used. Possible values:

- `ActiveDirectory`: Nested membership of all of a user's LDAP groups.
- `NoGroups`: No group support.
- `MemberDN`: MemberDN-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.

Configuring LDAP

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

Steps

1. Click **Cluster > LDAP**.
2. Click **Yes** to enable LDAP authentication.
3. Click **Add a Server**.
4. Enter the **Host Name/IP Address**.
5. Optional: Select **Use LDAPS Protocol**.
6. Enter the required information in **General Settings**.
7. Click **Enable LDAP**.
8. Click **Test User Authentication** if you want to test the server access for a user.
9. Optional: Click **Save Changes** to save any new settings.

Related tasks

[Creating a cluster admin account](#) on page 28

Related references

[LDAP details](#) on page 30

Disabling LDAP

You can disable LDAP integration using the Element OS web UI.

Before you begin

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

Steps

1. Click **Cluster > LDAP**.
2. Click **No**.
3. Click **Disable LDAP**.

Configuring cluster settings

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

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

Related concepts

[Using replication between SolidFire and ONTAP clusters](#) on page 107

[Working with virtual volumes](#) on page 68

[Managing Terms of Use](#) on page 34

[Managing SNMP](#) on page 35

[Managing drives](#) on page 37

[Managing nodes](#) on page 38

[Managing virtual networks](#) on page 42

Related tasks

[Setting cluster full threshold](#) on page 32

[Enabling and disabling encryption for a cluster](#) on page 33

[Enabling a broadcast client](#) on page 35

[Viewing Fibre Channel ports details](#) on page 41

Related information

[How to calculate SolidFire system error alert percentage](#)

Setting cluster full threshold

You can change the level at which the system generates a cluster fullness warning.

Before you begin

You have cluster admin 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 to calculate SolidFire system error alert percentage](#)

Encryption at rest

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

All drives in SF-series nodes 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 and must then be supplied for every read and write operation to the drive.

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 API or web UI, Encryption at Rest will be disabled on the drives. After the drive is removed, the drive can be secure 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.

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 admin 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.**

Managing Terms of Use

You can use the Element OS web UI to configure a banner that appears containing 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 OS web 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 admin 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 OS web UI.

Before you begin

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

Steps

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

Network time protocol

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 use the Element OS web UI to enter up to five different NTP servers.

Enabling a broadcast client

You can use the broadcast client setting to instruct each node in a cluster to listen for NTP broadcasts instead of querying an NTP server for updates.

Before you begin

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

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 use the Element OS web UI to configure Simple Network Management Protocol (SNMP) in your cluster. You can select which version of SNMP to use and configure traps to monitor the SolidFire cluster. You can also view and access management information base files.

SNMP details

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

SNMP MIBs

You can see which MIB files 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.
5. Optional: To add another requestor, follow these steps:
 - a. Click **Add a Requestor**.
 - b. 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.
6. Click **Save Changes**.

Related tasks

[Configuring SNMP traps](#) on page 36

[Viewing management information base files](#) on page 37

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, follow these steps:
 - a. Click **Add a USM User**.
 - b. Enter a 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, follow these steps:
 - a. Click **Add a Trap Recipient**.
 - b. Enter the host, port, and community string for the recipient.
 5. Click **Save Changes**.

Viewing 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 OS web UI to manage drives.

Related tasks

[Adding drives to a cluster](#) on page 21

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 table describes the elements shown 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.

Type

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

Wear remaining on drives

The storage system reports the approximate amount of wear available on each solid-state disk (SSD) drive 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.

Managing nodes

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

Related tasks

[Adding a node to a cluster](#) on page 38

Related references

[Node states](#) on page 40

Adding a node to a cluster

You can add nodes to a cluster when more storage is needed or during 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 OS 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.

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 OS than the version running on the cluster, the cluster asynchronously updates the node to the version of Element OS 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 39

[Configuring a Fibre Channel node](#) on page 19

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, and 38410 nodes can coexist in a cluster.

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.

Node versioning and compatibility

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

The following list describes the SolidFire software release significance levels that make up the Element OS 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 OS incremental releases with a different minor component. For example, 7.0 is not compatible with 7.1, and 7.1 is not compatible with 7.2.

Micro

The third number designates a compatible patch (incremental release) to the Element version represented by the major.minor components. For example, 7.0.1 is compatible with 7.0.2, and 7.0.2 is compatible with 7.0.3.

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

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.
- **PendingActive:** The system is in the process of installing compatible Element OS 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 system-generated node name.

Available 4k IOPS

Displays the IOPS configured for the node.

Node Role

Identifies what role the node has in the cluster. This can be Cluster Master, Ensemble Node, or Fibre Channel node.

Node Type

Displays 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 references

[Fibre Channel ports details](#) on page 41

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 43

[Enabling virtual routing and forwarding](#) on page 44

[Editing a virtual network](#) on page 44

[Editing VRF VLANs](#) on page 44

[Deleting a virtual network](#) on page 45

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

Shows if 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 SolidFire cluster.

Before you begin

- Identify the block of IP addresses that will be assigned to the virtual networks on the SolidFire nodes.
- Identify a storage network IP (SVIP) address that will be used as an endpoint for all SolidFire 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. You do not need to manually assign virtual network addresses to the nodes in the cluster.

Steps

1. Click **Cluster > Network**.
2. Click **Create VLAN**.
3. 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 43

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 44

Upgrading Element OS software

You must use the HealthTools suite of tools to upgrade the software starting with Element OS 10.2. This suite is made available by getting the latest version of the management node. The management node contains the software needed to perform the upgrade process.

Related concepts

[Using HealthTools for software upgrades \(Element OS 10.2 and later\)](#) on page 45

Related tasks

[Upgrading management node software](#) on page 49

[Upgrading Element OS 10.2 and later](#) on page 50

[Upgrading Element OS 10.2 and later on dark sites](#) on page 51

[Upgrading software for versions earlier than Element OS 10.2](#) on page 52

Using HealthTools for software upgrades (Element OS 10.2 and later)

The HealthTools suite includes the components required for the software upgrade process: sfupdate-healthtools, sfupgradecheck, and sfinstall. When an Element OS 10.2 version management node is either installed for the first time or upgraded from a previous version, HealthTools is included.

Note: The commands noted here require escalated privileges to run. Either preface commands with `sudo` or escalate your user to root privileges.

Related tasks

[Checking the installed version of HealthTools](#) on page 46

[Checking the installed version of HealthTools on dark sites](#) on page 46

[Updating HealthTools](#) on page 47

Checking the installed version of HealthTools

You can run the `sfupdate-healthtools` command to check the version of HealthTools installed on the management node. This will compare the currently installed version with the newest available version.

Before you begin

You are using a management node with Element OS 10.2.

Steps

1. To view the installed version, in the command line enter the following command:

```
sfupdate-healthtools -v
```

A sample of the output is as follows:

```
Currently installed version of HealthTools: 2018.02.01.100
```

2. To view the latest available version, in the command line enter the following command:

```
sfupdate-healthtools -l
```

A sample of the output is as follows:

```
Latest available version of HealthTools: 2018.02.01.200
The latest version of HealthTools can be downloaded from:
https://mysupport.netapp.com/NOW/cgi-bin/software/
```

Checking the installed version of HealthTools on dark sites

For management nodes isolated from external Internet access, the verification check for the latest version of the HealthTools always fails. You must manually verify if the latest version available is the same as the installed version.

Before you begin

- You have a management node on a dark site.
- When you run `sfupdate-healthtools -l`, you are unable to verify the latest available version of HealthTools.

Steps

1. Run the following command to check which version of HealthTools is installed on the management node:


```
sfupdate-healthtools -v
```
2. Go to https://library.netapp.com/ecm/ecm_get_file/ECMLP2840740 and check the `latest_version` field.

If your installed version of HealthTools is not the same as what is in the `latest_version` field, you must manually copy the newest HealthTools software to the management node using USB.
3. Log in to the NetApp Support Site at <https://mysupport.netapp.com/NOW/cgi-bin/software/> to download the latest HealthTools software.

Related tasks

[Updating HealthTools](#) on page 47

Updating HealthTools

You can use the `sfupdate-healthtools` command to update the version of HealthTools that is installed on the management node.

Steps

1. Log in to the NetApp Support Site at <https://mysupport.netapp.com/NOW/cgi-bin/software/>.
2. Download the latest version of HealthTools to a computer that is not the management node.
3. Copy the installation file to the management node.
You can do this using, for example, SCP.
4. Run the `sfupdate-healthtools <path to install file>` command to install the new software.

Example

See the following sample input command:

```
sfupdate-healthtools /tmp/solidfire-healthtools-2018.02.01.200.tgz
```

A sample of the output is as follows:

```
Checking key signature for file /tmp/solidfire-
healthtools-2018.02.01.200/components.tgz
installing command sfupdate-healthtools
Restarting on version 2018.02.01.200
sfupdate-healthtools /sf/bin/sfupdate-healthtools -r 2018.02.01.200
installing command sfupgradecheck
installing command sfinstall
installing command sfresetupgrade
```

5. Run the `sfupdate-healthtools -v` command to check the installed version.

Example

A sample of the output is as follows:

```
Currently installed version of HealthTools:
2018.02.01.200
```

6. If HealthTools is out-of-date, the upgrade check fails with an error message.

A sample output is as follows:

```
sfupgradecheck failed: HealthTools is out of date: installed version:
2018.02.01.200 latest version: 2018.03.05.901.
The latest version of the HealthTools can be downloaded from: https://
mysupport.netapp.com/NOW/cgi-bin/software/
Or rerun with the -n option
```

When this happens, you must either follow the steps described in this procedure or rerun `sfupgradecheck -n -u <cluster-user-name> -p <cluster-password> MVIP` to skip the version checking.

7. If your management node is on a dark site, the upgrade check fails.

A sample output is as follows:

```
sfupgradecheck failed: Unable to verify latest available version of
healthtools.
Manually verify latest version by checking 'latest_version' field in
https://library.netapp.com/ecm/ecm_get_file/ECMLP2840740 (this
version: 2018.02.01.200)
Re-run the command with '-n, --skip-version-check' parameter. Please
run 'sfupgradecheck --help' for full list of parameters.
```

When this happens, you must verify that the HealthTools are up-to-date and then run

```
sfupgradecheck -n -u <cluster-user-name> -p <cluster-password> MVIP or
sfupgradecheck --skip-version-check -u <cluster-user-name> -p <cluster-
password> MVIP.
```

Upgrade readiness checks

You can use the `sfupgradecheck` command to verify that the cluster is ready to be upgraded. This command verifies that there are no pending nodes, there is adequate disk space on the root partition, and each node can connect to the management node.

Scenarios

- If your cluster is upgrade ready, running the `sfupgradecheck -u <cluster-user-name> -p <cluster-password> MVIP` command does not give an error.

See the following sample input command:

```
sfupgradecheck -u admin -p admin 10.117.78.244
```

A sample output is as follows:

```
All node level tests passed.
Passed tests:
check_mnode_connectivity:
Test Description:    Verify storage nodes can communicate with
management node
Passed node IDs:    1, 2, 3, 4
More information:   https://kb.netapp.com/support/s/article/
kallA0000008ltYQAQ/mNode-connectivity
check_root_disk_space:
Test Description:    Verify node root directory has at least 12 GBs
of available disk space
Passed node IDs:    1, 2, 3, 4
More information:   https://kb.netapp.com/support/s/article/
kallA0000008ltTQAQ/SolidFire-Disk-space-error
check_pending_nodes:
Test Description:    Verify no pending nodes in cluster
More information:   https://kb.netapp.com/support/s/article/
kallA0000008ltOQAQ/pending-nodes
```

- If your cluster is not upgrade ready, checks within the `sfupgradecheck` tool will fail with error message details about the issue.

A sample output is as follows:

```
The following tests failed:
check_root_disk_space:
Severity:            ERROR
Test Description:    Verify node root directory has at least 12 GBs
of available disk space
Failed node IDs:    1
Remedy:              Remove unneeded files from root drive
More information:   https://kb.netapp.com/support/s/article/
```



```

kallA0000008ltTQAQ/SolidFire-Disk-space-error
Passed tests:
check_mnode_connectivity:
Test Description:      Verify storage nodes can communicate with
management node
Passed node IDs:      1, 2, 3, 4
More information:     https://kb.netapp.com/support/s/article/
kallA0000008ltYQAQ/mNode-connectivity
check_root_disk_space:
Test Description:      Verify node root directory has at least 12 GBs
of available disk space
Passed node IDs:      2, 3, 4
More information:     https://kb.netapp.com/support/s/article/
kallA0000008ltTQAQ/SolidFire-Disk-space-error
check_pending_nodes:
Test Description:      Verify no pending nodes in cluster
More information:     https://kb.netapp.com/support/s/article/
kallA0000008ltOQAQ/pending-nodes

```

In this example, node 1 is low on disk space. You can find more information in the knowledge base article listed in the error message.

Upgrading management node software

Before you upgrade the Element OS software on your cluster, you must ensure that you have the latest version of the management node software. To check the version of software running on your management node, look at the file `/solidfire/version` on your management node.

About this task

If you have a newly installed system, you must contact NetApp Support for information about whitelisting the public IP of the management node. The modification of the policy for `pubrepo.solidfire.com` enables you to fetch the repository and upgrade the management node.

Steps

1. Run the following script:

```
sudo update-fdva
```

Example

See the following sample output:

```

Begin time: Tue Jan  8 13:49:47 2013
[20]... [19]... [18]... [17]... [16]... [15]... [14]... [13]...
[12]... [11]... [10]... [9]... [8]... [7]... [6]... [5]... [4]...
[3]... [2]... [1]... [0]...
End time: Tue Jan  8 13:49:50 2013
Proceed indexes: [PPPPPP]

688.6 MiB will be downloaded into archive.
Downloading 913 archive files using 20 threads...
Begin time: Tue Jan  8 13:49:51 2013
[20]... [19]... [18]... [17]... [16]... [15]... [14]... [13]...
[12]... [11]... [10]... [9]... [8]... [7]... [6]... [5]... [4]...
[3]... [2]... [1]... [0]...
End time: Tue Jan  8 13:54:36 2013
0.0 bytes in 0 files and 0 directories can be freed.
Run /solidfire-ubuntu-mirror/var/clean.sh for this purpose.
Running the Post Mirror script ...
(/solidfire-ubuntu-mirror/var/postmirror.sh)
Post Mirror script has completed. See above output for any possible
errors.

>>> Mirroring GCC repository

```

```

Downloading 7 index files using 7 threads...
Begin time: Tue Jan  8 13:58:13 2013
[7]... [6]... [5]... [4]... [3]... [2]... [1]... [0]...
End time: Tue Jan  8 13:58:13 2013
Proceed indexes: [P]
21.2 MiB will be downloaded into archive.
Downloading 11 archive files using 11 threads...
Begin time: Tue Jan  8 13:58:13 2013
[11]... [10]... [9]... [8]... [7]... [6]... [5]... [4]... [3]...
[2]... [1]... [0]...
End time: Tue Jan  8 13:58:21 2013
0.0 bytes in 0 files and 0 directories can be freed.
Run /solidfire-ubuntu-mirror/var/clean.sh for this purpose.
Running the Post Mirror script ...
(/solidfire-ubuntu-mirror/var/postmirror.sh)
Post Mirror script has completed. See above output for any possible
errors.

>>> Mirroring OMSA repository
Downloading 4 index files using 4 threads...
Begin time: Tue Jan  8 13:58:21 2013
[4]... [3]... [2]... [1]... [0]...
End time: Tue Jan  8 13:58:21 2013
Proceed indexes: [P]
197.2 MiB will be downloaded into archive.
Downloading 73 archive files using 20 threads...
Begin time: Tue Jan  8 13:58:21 2013
[20]... [19]... [18]... [17]... [16]... [15]... [14]... [13]...
[12]... [11]... [10]... [9]... [8]... [7]... [6]... [5]... [4]...
[3]... [2]... [1]... [0]...
End time: Tue Jan  8 13:59:24 2013
0.0 bytes in 0 files and 0 directories can be freed.
Run /solidfire-ubuntu-mirror/var/clean.sh for this purpose.
Running the Post Mirror script ...
(/solidfire-ubuntu-mirror/var/postmirror.sh)
Post Mirror script has completed. See above output for any possible
errors.

>>> Updating local APT repository cache
done

>>> Installing latest version of sfinstall
Reading package lists... Done
Building dependency tree
Reading state information... Done

```

2. Restart the management node.

Upgrading Element OS 10.2 and later

You must use the `sfinstall` file included in the HealthTools suite of tools to update the Element software for version 10.2 and later.

Before you begin

You have the latest version of the management node software.

Steps

1. Log in to the NetApp Support Site.
2. Download the newest version of Element OS to a computer that is not the management node.
3. Copy the ISO file to the management node in an accessible location like `/tmp`.
You can do this using, for example, SCP.
4. Run the `sfinstall` command with the following parameters to give the path to the ISO file:

```
sfinstall <MVIP> -u <cluster_username> -p <cluster_password> <path-to-install-file-ISO>
```

Example

See the following sample input command:

```
sfinstall 10.117.78.244 -u admin -p admin /tmp/solidfire-rtfi-neon-patch2.10.2.0.34.iso
```

The output for the sample shows that `sfinstall` attempts to verify if a newer version of `sfinstall` is available:

```
sfinstall 10.117.78.244 -u admin -p admin /tmp/solidfire-rtfi-neon-patch2-10.2.0.34.iso
2018-01-24 16:52:15: Newer version of sfinstall available. This version: 2018.02.01.200, latest version: 2018.06.05.901. The latest version of the HealthTools can be downloaded from: https://mysupport.netapp.com/NOW/cgi-bin/software/ or rerun with --skip-version-check
```

5. To list available versions of software, use the `-l` option:

Example

See the following sample input command:

```
sfinstall 10.117.78.148 -l -u admin -p admin
```

The output for the sample is as follows:

```
2018-01-24 16:35:37: sfinstall Release Version: 2018.02.01.200
Revision: 09aafd6a897d Build date: 2018-01-23 13:51
2018-01-24 16:35:37: Checking connectivity to MVIP 10.117.78.148
2018-01-24 16:35:37: Successfully connected to cluster MVIP
2018-01-24 16:35:38: Currently installed:
2018-01-24 16:35:38:   solidfire-rtfi-neon-patch2-10.2.0.87 10.2.0.87
2018-01-24 16:35:38: Available to install:
2018-01-24 16:35:38:   10.2.0.99                                10.2.0.99
```

Related tasks

[Upgrading management node software](#) on page 49

Upgrading Element OS 10.2 and later on dark sites

You must use the `sfinstall` file included in the HealthTools suite of tools to update the Element software for version 10.2 and later.

About this task

- You have the latest version of the management node software.
- Your management node is not connected to the Internet.

Steps

1. Log in to the NetApp Support Site.
2. Download the newest version of Element OS to a computer that is not the management node.

3. Copy the ISO file to the management node in an accessible location like /tmp.

You can do this using, for example, SCP.

4. Run the `sfinstall` command with the following parameters to give the path to the ISO file:

```
sfinstall <MVIP> -u <cluster_username> --skip-version-check -p  
<cluster_password> <path-to-install-file-ISO>
```

When this is run on a dark site, the command fails unless you run the command with the `--skip-version-check` option:

See the following sample input command:

```
sfinstall 10.117.78.244 -u admin -p admin /tmp/solidfire-rtfi-neon-  
patch2.10.2.0.34.iso.
```

The output for the sample is as follows:

```
sfinstall failed: Unable to verify latest available version of  
healthtools.  
Manually verify latest version by checking 'latest_version' field in  
https://library.netapp.com/ecm/ecm_get_file/ECMLP2840740 (this  
version: 2018.02.01.200)  
Re-run the command with '-n, --skip-version-check' parameter.
```

Related tasks

[Checking the installed version of HealthTools on dark sites](#) on page 46

[Upgrading management node software](#) on page 49

Upgrading software for versions earlier than Element OS 10.2

You must update the Element software on the nodes from one release to the next. You might have to install patch updates when they become available. The software installation files are located on the management node.

Before you begin

You have the latest version of the management node software.

Steps

1. Run the `sfinstall` command with the following parameters to set the repository to point to the new management node:

```
sfinstall <MVIP> -u <cluster_username> -p <cluster_password> -s <mNode  
IP address> <name of repository>
```

Example

See the following sample input command:

```
sfinstall 192.168.143.113 -u admin -p solidfire -s 192.168.128.100  
"mnode repository"
```

The output for the sample is as follows:

```
2012-06-28 10:40:01: Repositories:  
2012-06-28 10:40:01: 192.168.128.100 - mnode repository
```

2. Run the `sfindinstall` command with the following parameters to get the name of the software package to install:

```
sfindinstall <MVIP> -u <cluster_username> -p <cluster_password> -l
```

Example

See the following sample input command:

```
sfindinstall 192.168.143.113 -u admin -p solidfire -l
```

The output for the sample is as follows:

```
Currently installed:
solidfire-san-nitrogen-7.nn.nn          7.nn.nn
Available to install:
solidfire-san-oxygen-8.nn.nn           8.nn.nn
solidfire-san-oxygen-patch1-8.nn.nn    8.nn.nn
```

3. Run the `sfindinstall` command with the following parameters to install the software:

```
sfindinstall <MVIP> -u <cluster_username> -p <cluster_password> solidfire-
san-[element]-[version number]
```

Example

See the following sample input command:

```
sfindinstall 192.168.143.113 -u admin -p solidfire solidfire-san-
[element]-[version number]
```

The output for the sample is as follows:

```
2012-06-28 12:54:07: Install of solidfire-san-[element]-[version
number] complete.
```

Related tasks

[Upgrading management node software](#) on page 49

Data management

You can manage data in your SolidFire cluster from the Management tab in the Element OS web 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 accounts](#) on page 54

[Working with volumes](#) on page 56

[Working with virtual volumes](#) on page 68

[Working with access groups and initiators](#) on page 77

Working with accounts

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

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 54

[Viewing individual account details](#) on page 55

[Editing an account](#) on page 55

[Deleting an account](#) on page 56

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 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.

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.

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

[SolidFire Quality of Service](#) on page 57

Related tasks

[Creating a QoS policy](#) on page 59

- [Editing a QoS policy](#) on page 60
- [Deleting a QoS policy](#) on page 60
- [Creating a volume](#) on page 61
- [Viewing individual volume performance details](#) on page 62
- [Editing active volumes](#) on page 63
- [Deleting a volume](#) on page 64
- [Restoring a deleted volume](#) on page 65
- [Purging a volume](#) on page 65
- [Cloning a volume](#) on page 65
- [Assigning LUNs to Fibre Channel volumes](#) on page 67
- [Applying a QoS policy to volumes](#) on page 67
- [Removing the QoS policy association of a volume](#) on page 67

SolidFire Quality of Service

A SolidFire 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.

IOPS parameters are defined in the following ways:

Minimum IOPS

The minimum number of sustained inputs and outputs per second (IOPS) that the SolidFire 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 SolidFire 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.

SolidFire 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			
Parameters	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
<p>*These estimations are approximate.</p> <p>**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.</p>						

QoS performance curve

The SolidFire 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. SolidFire software takes into account the block sizes it receives by normalizing block sizes to 4k. Based on workload, the system may 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.

Related tasks

[Creating a QoS policy](#) on page 59

[Editing a QoS policy](#) on page 60

[Deleting a QoS policy](#) on page 60

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.

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.

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. Names can be up to 223 characters in length and contain a-z, 0-9, and dash (-) characters.

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

Identifies if 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.
 - The default volume size selection is in GB. You can adjust volume 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: 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 a SolidFire 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.

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 SolidFire 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.

Caution: 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 59

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**.

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 OS web UI.

The SolidFire 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 OS web 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 SolidFire APIs for the SolidFire storage system. Commands to create, delete, and manage storage containers and delete virtual volumes can be initiated using the Element OS web UI.

The majority of configurations necessary for using Virtual Volumes functionality with SolidFire 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: The VVol functionality has a cluster limit of 8000 virtual volumes.

Related concepts

[Protocol endpoints](#) on page 75

[Bindings](#) on page 76

Related tasks

[Enabling virtual volumes](#) on page 68

[Deleting a virtual volume](#) on page 72

[Creating a storage container](#) on page 73

[Editing a storage container](#) on page 74

[Deleting a storage container](#) on page 74

Enabling virtual volumes

You must manually enable vSphere Virtual Volumes (VVols) functionality through the SolidFire Element OS web UI. The SolidFire 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 SolidFire cluster must be running Element OS version 9.0 or later.
- The SolidFire cluster must be connected to an ESXi 6.0 and later environment that is compatible with VVols.

About this task

Attention: Enabling vSphere Virtual Volumes functionality permanently changes Element OS 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.

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 OS web 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. Go to **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 OS user interface. 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 OS web UI for the cluster.

- You have created an associated storage container.
- You have configured your vSphere cluster to use SolidFire 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 OS web UI. You should only delete a virtual volume from the Element OS web 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 OS web 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 OS web 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 SolidFire cluster.

Storage containers are created and tied to accounts on SolidFire. A storage container created on SolidFire storage appears as a vSphere datastore in vCenter and ESXi. Storage containers do not allocate any space on SolidFire 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 SolidFire Element OS web 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 OS web 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 SolidFire 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 SolidFire 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

Indicates the protocol endpoint type (SCSI is the only available protocol for Element OS).

Storage Container ID

The UUID of the virtual volume storage container.

Active 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 SolidFire 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 OS web 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 OS web 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 SolidFire cluster. 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 SolidFire cluster automatically creates one protocol endpoint per storage node in the cluster. For example, a six-node SolidFire cluster has six protocol endpoints that are mapped to each ESXi host. Protocol endpoints are dynamically managed by SolidFire 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 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

Indicates 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 access groups and initiators

You can use iSCSI initiators or Fibre Channel initiators to access 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. 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 an access group](#) on page 77

[Adding volumes to an access group](#) on page 79

[Removing volumes from an access group](#) on page 79

[Creating an initiator](#) on page 80

[Editing an initiator](#) on page 81

[Adding a single initiator to an access group](#) on page 81

[Adding multiple initiators to an access group](#) on page 82

[Removing initiators from an access group](#) on page 82

[Deleting an access group](#) on page 83

[Deleting an initiator](#) on page 83

Creating an 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.

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	<p>a. Under Add Initiators, select an existing Fibre Channel initiator from the Unbound Fibre Channel Initiators list.</p> <p>b. Click Add FC Initiator.</p> <p>Note:</p> <p>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.</p> <p>A sample of the format is as follows:</p> <pre>5f:47:ac:c0:5c:74:d4:02</pre>
Adding an iSCSI initiator	<p>Under Add Initiators, select an existing initiator from the Initiators list.</p> <p>Note:</p> <p>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.</p> <p>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>

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 79

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.

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.

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 an access group

You can add an initiator to an existing 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.

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 an access group

You can add multiple initiators to an existing volume access group to allow access to volumes in the volume access group 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.

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

Element OS enables you to protect your data in a variety of ways with capabilities such as snapshots for individual volumes or groups of volumes, snapshot scheduling capability, replication between clusters and volumes, and replication to ONTAP systems.

Related concepts

- [Using individual volume snapshots](#) on page 84
- [Using group snapshots](#) on page 89
- [Using snapshot schedules](#) on page 94
- [Using replication between SolidFire clusters](#) on page 96
- [Using replication between SolidFire and ONTAP clusters](#) on page 107
- [Backing up and restoring volumes](#) on page 120

Using individual volume snapshots

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.

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 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; you can also create a clone of a volume from a replicated snapshot.

Related tasks

- [Creating a volume snapshot](#) on page 84
- [Editing snapshot retention](#) on page 85
- [Deleting a snapshot](#) on page 86
- [Cloning a volume from a snapshot](#) on page 86
- [Rolling back a volume to a snapshot](#) on page 87
- [Backing up a volume snapshot to an Amazon S3 object store](#) on page 87
- [Backing up a volume snapshot to an OpenStack Swift object store](#) on page 88
- [Backing up a volume snapshot to a SolidFire cluster](#) on page 88

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 concepts

[Using snapshot schedules](#) on page 94

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 SolidFire cluster. 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

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 certain state.

Related tasks

[Creating a group snapshot](#) on page 90

[Editing group snapshots](#) on page 91

[Editing members of group snapshot](#) on page 92

[Deleting a group snapshot](#) on page 91

[Rolling back volumes to a group snapshot](#) on page 92

[Cloning multiple volumes](#) on page 93

[Cloning multiple volumes from a group snapshot](#) on page 93

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

Identifies 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**.

Using snapshot schedules

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.

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 94

[Editing a snapshot schedule](#) on page 95

[Deleting a snapshot schedule](#) on page 96

[Copying a snapshot schedule](#) on page 96

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**.

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.

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**.

Using replication between SolidFire clusters

NetApp SolidFire Real-Time Replication technology enables the quick creation of remote copies of volume data. You can pair a SolidFire cluster with up to four other clusters. You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair for failover and failback scenarios.

Related tasks

[Pairing clusters](#) on page 97

[Pairing volumes](#) on page 100

Configuring cluster and volume pairing for real-time replication

You must first pair two SolidFire clusters and then pair volumes on each cluster to take advantage of real-time remote replication.

Before you begin

- Cluster admin 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.

- The difference between Element OS 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.

Steps

1. Create a cluster pair.
2. Create a volume pair.
3. Assign a replication source and target to the paired volumes.

Related tasks

[Pairing clusters](#) on page 97

[Pairing volumes](#) on page 100

[Assigning a replication source and target to paired volumes](#) on page 103

Cluster pairs

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).

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.

Pairing clusters

You can pair a source and target cluster using the MVIP of the target cluster if there is Cluster Admin access to both clusters. If Cluster Admin 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.

Before you begin

- Cluster admin 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.
- The difference between Element OS 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.

Step

1. Select one of the following methods to pair clusters:
 - Pair clusters using MVIP: Use this method if there is Cluster Admin 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 Admin 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 98

[Pairing clusters using a pairing key](#) on page 98

Related references

[Network port requirements](#) on page 11

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 Admin access on both of clusters is required to use this method. The Cluster Admin 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 admin 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.

In the **Cluster Pairs** window, verify that the cluster pair is connected.

(Optional) Navigate back to the local cluster and the **Cluster Pairs** window and verify that the cluster pair is connected.

Related tasks

[Pairing clusters using a pairing key](#) on page 98

Related information

[Pairing clusters using MVIP \(video\)](#)

Pairing clusters using a pairing key

If you have Cluster Admin 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 Admin 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 Admin 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 the **Complete Pairing** button.

In the **Cluster Pairs** window, verify that the cluster pair is connected.

(Optional) Navigate back to the local cluster and the **Cluster Pairs** window and verify that the cluster pair is connected.

Related tasks

[Pairing clusters using MVIP](#) on page 98

Related information

[Pairing clusters using a cluster pairing key \(video\)](#)

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.

Deleting a cluster pair

You can delete a cluster pair from the Element OS web 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.

Volume pairs

You can pair two volumes for real-time replication that are stored on different 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.

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 admin privileges to one or both clusters being paired.

Steps

1. Pair volumes by selecting one of the following methods:
 - Using a volume ID: Use this method if you have cluster admin 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 admin 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.

2. Assign a replication target and source to the paired volumes.

Note: A replication source volume has read/write account access. A replication target volume can only be accessed by the replication source as read/write.

3. Optional: Verify that the cluster pairs have an active connection.

Related tasks

[Pairing clusters using MVIP](#) on page 98

[Pairing clusters using a pairing key](#) on page 98

[Pairing volumes using a volume ID](#) on page 101

[Pairing volumes using a pairing key](#) on page 102

[Assigning a replication source and target to paired volumes](#) on page 103

Related references

[Network port requirements](#) on page 11

Pairing volumes using a volume ID

You can pair a volume with another volume on a remote cluster if you have cluster admin 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, 512e setting, 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.

About this task

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, 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 web UI of the remote cluster. You might be required to log on to the remote cluster with cluster admin credentials.

10. In the web 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 103

Related references

[Volume pairing messages](#) on page 105

[Volume pairing warnings](#) on page 106

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, 512e setting, 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 web 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 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 103

Related references

[Volume pairing messages](#) on page 105

[Volume pairing warnings](#) on page 106

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. From the cluster that contains the volume you want to assign as source, select **Management > Volumes**.
2. Click the Actions icon for the volume you want to assign as source.
3. Click **Edit**.
4. 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.

5. Click **Save Changes**.
6. From the cluster that contains the volume you want to assign as target, select **Management > Volumes**.
7. Click the Actions icon for the volume you want to assign as target.
8. Click **Edit**.
9. 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.
10. Click **Save Changes**.
11. Optional: From either cluster, select **Data Protection > Volume Pairs**.
12. Optional: Verify that the volume status is `Active`.

Related tasks

[Pairing volumes using a volume ID](#) on page 101

[Pairing volumes using a pairing key](#) on page 102

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

Indicates the client write replication method. Possible values are as follows:

- Async
- Snapshot-Only
- Sync

Direction

Indicates 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.

Paused Disconnected

Source replication or sync RPCs timed out. Connection to the remote cluster has been lost. Check network connections to the cluster.

Resuming Connected

The remote replication sync is now active. Beginning the sync process and waiting for data.

Resuming RR Sync

A single helix copy of the volume metadata is being made to the paired cluster.

Resuming Local Sync

A double helix copy of the volume metadata is being made to the paired cluster.

Resuming Data Transfer

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.

Paused Cluster Full

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.

Paused Manual

Local volume has been manually paused. It must be unpaused before replication resumes.

Paused Manual Remote

Remote volume is in manual paused mode. Manual intervention required to unpause the remote volume before replication resumes.

Paused Misconfigured

Waiting for an active source and target. Manual intervention required to resume replication.

Paused QoS

Target QoS could not sustain incoming IO. Replication auto-resumes. The message displays on the source end of the pair only.

Paused Slow Link

Slow link detected and stopped replication. Replication auto-resumes. The message displays on the source end of the pair only.

Paused Volume Size Mismatch

Target volume is smaller than 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.

Stopped Misconfigured

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.

Editing volume pairs

You can edit volume pair properties to make changes to the replication mode of the volume pair relationship or manually pause replication.

Steps

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon that corresponds to the volume pair you want to edit.
3. Click **Edit**.
4. In the **Edit Volume Pair** pane, do any of the following:
 - **Attention:** Pausing or resuming volume replication manually will cause the transmission of data to cease or resume. Changing the mode of replication will cause the mode to change immediately. Be sure that you have coordinated these changes at both sites.
 - Manually pause or start the replication process.
 - 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.
5. Click **Save Changes**.

Deleting volume pairs

You can delete a volume pair if wish to remove a pair association between two volumes.

Steps

1. Select **Data Protection > Volume Pairs**.
2. Click the Actions icon that corresponds to the volume pair you want to delete.
3. Click **Delete**.
4. Confirm the message.

Using replication between SolidFire and ONTAP clusters

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

Related concepts

- [SnapMirror endpoints](#) on page 109
- [SnapMirror labels](#) on page 111
- [SnapMirror relationships](#) on page 112
- [Disaster recovery using SnapMirror](#) on page 116

Related tasks

[Enabling SnapMirror on the cluster](#) on page 108

[Enabling SnapMirror on the volume](#) on page 109

SnapMirror overview

Beginning with Element OS 10.1, SolidFire supports SnapMirror functionality to copy and restore snapshots with ONTAP systems.

Element OS 10.1 and above includes code that can communicate directly with SnapMirror on ONTAP systems running 9.3 or higher. The SolidFire API provides methods to enable SnapMirror functionality on clusters, volumes, and snapshots. Additionally, the SolidFire Element OS web UI includes all necessary functionality to manage SnapMirror relationships between Element OS and ONTAP based systems.

Starting with Element OS 10.3 and ONTAP 9.4 systems, you can replicate ONTAP volumes to SolidFire in specific use cases with limited functionality. For more information, see ONTAP documentation.

Related information

[Replication between SolidFire Element OS and ONTAP](#)

Enabling SnapMirror on the cluster

You must manually enable SnapMirror functionality at the cluster level through the SolidFire Element OS web UI. The SolidFire 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 SolidFire cluster must be running Element OS version 10.1 or later.

About this task

SnapMirror can only be enabled for SolidFire clusters 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 Element OS 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. This allows replication of SolidFire data to specified ONTAP volumes. This is permission from the SolidFire cluster administrator for SnapMirror to control a volume.

Before you begin

- You have enabled SnapMirror in the Element OS web UI for the cluster.
- A SnapMirror endpoint is available.
- The volume must be 512e block size.
- The volume is not participating in SolidFire 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**.

SnapMirror endpoints

A SnapMirror endpoint is an ONTAP cluster that serves as a replication target for a SolidFire cluster. You must create a SnapMirror endpoint before you can create a SnapMirror relationship.

You can create and manage up to four SnapMirror endpoints on a SolidFire cluster.

Note: If an existing endpoint was originally created using the API and credentials were not saved, you can see the endpoint in the UI and verify its existence, but it cannot be managed using the Element OS web UI. This endpoint can then only be managed using the API.

Creating an endpoint

You must create a SnapMirror endpoint in the SolidFire Element OS web UI before you can create a relationship.

Before you begin

- You have enabled SnapMirror in the Element OS web UI for the cluster.
- You know the ONTAP credentials for the endpoint.

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 admin credentials associated with the endpoint.
5. Click **Create Endpoint**.

SnapMirror endpoint details

The SnapMirror Endpoints page on the Data Protection tab provides information about all SnapMirror endpoints on the cluster such as ID, cluster name, and cluster management IP.

ID

The ID of the endpoint.

Cluster Name

Name of the destination cluster.

Cluster Management IP

IP address of the destination cluster.

LIFs

This lists the ONTAP intercluster logical interfaces used to communicate with SolidFire.

Relationships

Number of relationships associated with this endpoint.

Status

Current status of the SnapMirror endpoint. Possible values are as follows:

- connected
- disconnected
- unmanaged

Editing an endpoint

You must modify a SnapMirror endpoint in the SolidFire Element OS web UI.

Before you begin

- You have enabled SnapMirror in the Element OS web UI for the cluster.
- An existing SnapMirror endpoint is available to modify.

Steps

1. Click **Data Protection > SnapMirror Endpoints**.
2. Click the Actions icon for the endpoint you want to edit.
3. In the resulting menu, select **Edit**.
4. Under **Cluster Management IP**, edit the IP address as needed.
5. Under **ONTAP Credentials**, edit the user name or password, if necessary.
6. Click **Save Changes**.

Deleting an endpoint

You must delete a SnapMirror endpoint from the SolidFire Element OS web UI.

Before you begin

- You have enabled SnapMirror in the Element OS web UI for the cluster.
- An existing SnapMirror endpoint is available to delete.

Steps

1. Click **Data Protection > SnapMirror Endpoints**.
2. Click the Actions icon for the endpoint you want to delete.
3. In the resulting menu, select **Delete**.
4. Confirm the action.
5. Refresh the list of SnapMirror endpoints to confirm that the endpoint has been removed.

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 OnCommand 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 OnCommand 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 94

SnapMirror relationships

A SnapMirror relationship is the relationship between a source volume and a destination volume. Data is replicated to the destination volume by using NetApp snapshot copies. You can edit and delete SnapMirror relationships using the SolidFire Element OS web UI.

Creating a SnapMirror relationship

You must create a SnapMirror relationship in the SolidFire Element OS web 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 web 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 OS web 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. 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.
12. Click **Create Relationship**.
13. Click **Data Protection > SnapMirror Relationships** to view this new SnapMirror relationship.

SnapMirror relationship details

The SnapMirror Relationships page on the Data Protection tab provides information about all SnapMirror relationships on the cluster such as endpoint ID, name of destination cluster, and name of destination volume.

Endpoint ID

The ID of the endpoint.

Source Cluster

Name of the source cluster.

Source Volume

Name of the source volume.

Destination Cluster

Name of the destination ONTAP cluster.

Destination Volume

Name of the destination ONTAP volume.

State

Current relationship state of the destination volume. Possible values are as follows:

- 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 as follows:

- async_mirror
- mirror_vault

Schedule Name

Name of the pre-existing schedule on the ONTAP system selected for this relationship.

Editing a SnapMirror relationship

You must edit a SnapMirror relationship in the SolidFire Element OS web UI.

Before you begin

- SnapMirror is enabled on the volume.
- An existing SnapMirror relationship is available to modify.

Steps

1. Click **Data Protection > SnapMirror Relationships**.

2. Click the Actions icon for the relationship that is to be edited.
3. Click **Edit**.
4. In the **Edit SnapMirror Relationship** dialog box, you can change the policy, schedule, or bandwidth limit settings.
5. Click **Save Changes**.

Deleting a SnapMirror relationship

You can delete a SnapMirror relationship from the Element OS web UI.

Before you begin

- You have enabled SnapMirror in the Element OS web UI for the cluster.
- An existing SnapMirror relationship is available to delete.

Steps

1. Click **Data Protection > SnapMirror Relationships**.
2. Click the Actions icon for the relationship you want to delete.
3. In the resulting menu, select **Delete**.
4. Confirm the action.
5. Refresh the list of SnapMirror relationships to confirm that the relationship has been removed.

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.

Disaster recovery using SnapMirror

In the event of a problem with a SolidFire volume or cluster, 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

You can perform a failover from the SolidFire cluster to make the destination volume read/write and accessible to hosts on the destination side.

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 SolidFire Element OS web UI to perform the failover. If the Element OS web UI is not available, you can also use OnCommand System Manager or ONTAP CLI to issue the break relationship command.

Steps

1. In the Element OS web 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.

SnapMirror failback to SolidFire

When the issue on the primary side has been mitigated, you must resynchronize the original source volume and fail back. 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 119

Related tasks

[Performing a failback when source volume still exists](#) on page 117

[Performing a failback when source volume no longer exists](#) on page 118

Performing a failback when source volume still exists

You can resynchronize the original source volume and fail back using the SolidFire Element OS web UI. This procedure applies to scenarios where the original source volume still exists.

Steps

1. In the Element OS web 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 failback scenarios](#) on page 119

Performing a failback when source volume no longer exists

You can resynchronize the original source volume and fail back using the SolidFire Element OS user interface. 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 SolidFire and ONTAP volumes.
- The SolidFire volume is irretrievably lost.
- The original volume name shows as NOT FOUND.

Steps

1. In the Element OS web 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 volume on SolidFire 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 failback scenarios](#) on page 119

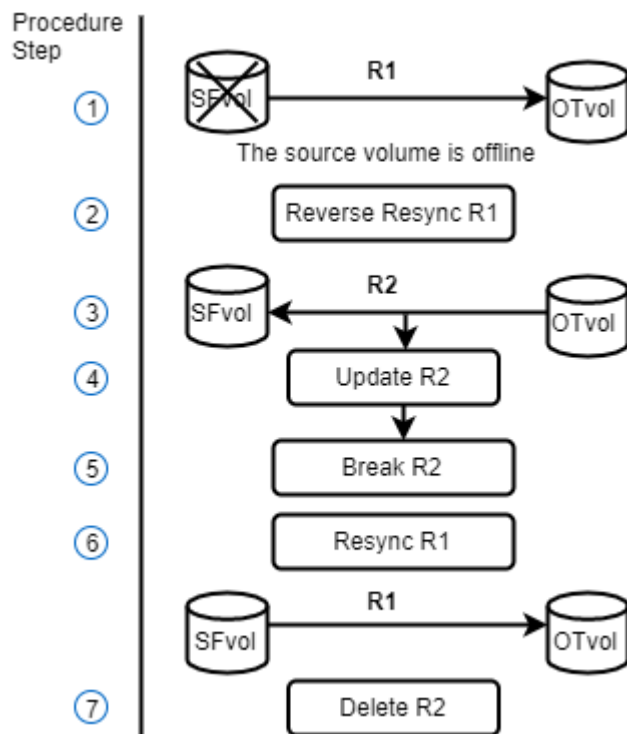
SnapMirror failback scenarios

The SnapMirror disaster recovery functionality is illustrated in two failback scenarios. These assume the original relationship has been failed over (broken).

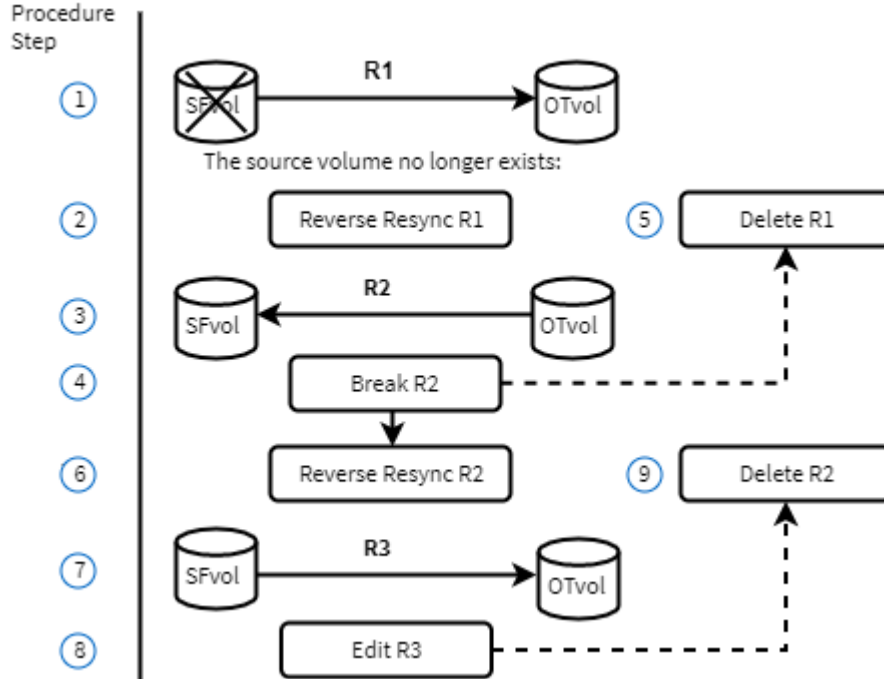
The steps from the corresponding procedures are added for reference.

Note: In the examples here, R1 = the original relationship in which SolidFire is the original source volume (SFvol) and ONTAP is the original destination volume (OTvol). 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 117

[Performing a failback when source volume no longer exists](#) on page 118

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.

When you restore volumes from OpenStack Swift or Amazon S3, you need manifest information from the original backup process. If you are restoring a SolidFire 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 120

[Backing up a volume to an OpenStack Swift object store](#) on page 121

[Backing up a volume to a SolidFire cluster](#) on page 121

[Restoring a volume from backup on an Amazon S3 object store](#) on page 122

[Restoring a volume from backup on an OpenStack Swift object store](#) on page 123

[Restoring a volume from backup on a SolidFire cluster](#) on page 124

Backing up a volume to an Amazon S3 object store

You can back up SolidFire 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 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 SolidFire 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 cluster

You can back up volumes residing on a SolidFire cluster to a remote SolidFire cluster.

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**.

Related concepts

[Cluster pairs](#) on page 97

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 cluster

You can restore a volume from a backup on a SolidFire 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 solid-state disk (SSD) drives for maintenance purposes.

Related concepts

- [Troubleshooting drives](#) on page 134
- [Troubleshooting nodes](#) on page 137
- [Working with per-node utilities](#) on page 138
- [Working with the management node](#) on page 142
- [Understanding cluster fullness levels](#) on page 146

Related tasks

- [Viewing information about system events](#) on page 125
- [Viewing system alerts](#) on page 128
- [Viewing node performance activity](#) on page 130
- [Viewing volume performance](#) on page 131
- [Viewing iSCSI sessions](#) on page 132
- [Viewing Fibre Channel sessions](#) on page 133
- [Enabling FIPS 140-2 on your cluster](#) on page 147
- [Enabling and disabling support access](#) on page 149

Related references

- [Viewing status of running tasks](#) on page 128

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

1. In the Element OS web 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).

Item	Description
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 references

[Event types](#) on page 126

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.

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.

incorrectBondPortCount

Events related to the number of Bond10G interfaces detected on the node.

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.

statEvent

Events related to system statistics.

sliceEvent

Events related to the Slice Server, such as removing a metadata drive or volume.

snmpTrapEvent

Events related to SNMP traps.

tsEvent

Events related to the system transport service.

unexpectedException

Events related to unexpected system exceptions.

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 OS web 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 OS web 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).

Item	Description
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 references

[Alert error codes](#) on page 129

Related information

[SolidFire Element OS API Reference Guide](#)

Alert error codes

The system reports error codes with each alert on the Alerts page. Error codes help you determine what component of the system experienced the alert, and you can learn more about why the alert was generated using the information in the Details column.

The following list outlines the different types of system alerts:

blockServiceTooFull

A block service is using too much space and running low on capacity.

blockServiceUnhealthy

The SolidFire Application cannot communicate with a Block Service. If this condition persists, the system relocates the data to another drive. After the system relocates the data, you should reboot the unhealthy node to restore communication.

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.

clusterFull

Stage 3 Cluster Full: Add additional capacity or free up capacity as soon as possible.
 Stage 4 Cluster Full: Due to high capacity consumption Helix data protection will not recover if a node fails. Creating new Volumes or Snapshots is not permitted until additional capacity is available. Add additional capacity or free up capacity immediately.

clusterIOPSAreOverProvisioned

The sum of all minimum QoS IOPS is greater than the expected IOPS of the cluster. The system cannot maintain minimum QoS in this condition. IOPS might need to be adjusted.

disconnectedClusterPair

Paired clusters have become disconnected. Reestablish communication between the clusters.

driveWearFault

A drive might need attention due to wear.

ensembleDegraded

Power or network connectivity has been lost to one or more of the ensemble nodes. Restore network connectivity or power to the affected node.

exception

A non-routine fault has been detected. This fault will not be cleared. Call NetApp Support to resolve the exception fault.

failedSpaceTooFull

A Slice Service is using space reserved for failed writes. Contact NetApp Support.

fibreChannelAccessDegraded

A node has stopped responding to the storage nodes in the cluster.

fibreChannelAccessUnavailable

All Fibre Channel nodes have become disconnected.

inconsistentMtus

Bond1G mismatch: Inconsistent MTUs detected on Bond1G interfaces. MTU to nodeID mapping: <mapping of MTUs to nodes>. Bond10G mismatch: Inconsistent MTUs detected on Bond10G interfaces. MTU to nodeID mapping: <mapping of MTUs to nodes>.

invalidConfiguredFibreChannelNodeCount

There is only one Fibre Channel node configured in a cluster. For proper Fibre Channel operation, at least two Fibre Channel nodes must be configured in a cluster.

nodeHardwareFault

The system has detected a hardware misconfiguration or a component that is not functioning as expected.

notUsingLACPBondMode

LACP bonding mode is not configured. We recommend using LACP bonding when deploying SF-Series 19210 and newer nodes; clients might experience timeouts if LACP is not configured.

provisionedSpaceTooFull

The overall provisioned capacity of the cluster is too full.

sliceServiceUnhealthy

The SolidFire Application cannot communicate with a metadata service.

sliceServiceTooFull

A Slice Service is using too much space and running low on capacity.

upgrade

The software on one or more nodes is being upgraded.

unbalancedMixedNodes

The storage on the mix of nodes in a cluster has become unbalanced in a way that might degrade performance.

volumeDegraded

Secondary volumes have not finished replicating and synchronizing.

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 OS web 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 references

[Volume performance details](#) on page 131

Volume performance details

You can view performance statistics of volumes from the Volume Performance page of the Reporting tab in the Element OS web 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. Values:

- 0 = Client is not using the volume
- 100 = Client is using their max
- >100 = Client is using their 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 OS web UI, select **Reporting > iSCSI Sessions**.
2. To see the filter criteria fields, click **Filter**.

Related references

[iSCSI session details](#) on page 133

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 OS web UI, select **Reporting > FC Sessions**.
2. To see the filter criteria fields, click **Filter**.

Related references

[Fibre Channel session details](#) on page 134

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 disk (SSD) drive with a replacement drive. SolidFire SSDs 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.

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 OS 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 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 135

Related tasks

[Removing failed drives](#) on page 135

[Removing MDSS drives](#) on page 137

[Replacing failed drives in the chassis](#) on page 137

Removing failed drives

The 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 5.5 minutes or longer. The system displays a list of the failed drives. You can remove a failed drive from the failed drive list.

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 5.5 minutes, the drives automatically update and continue as active drives in the cluster.

Steps

1. Click **Cluster > Drives**.
2. Click **Failed** to view the list of failed drives.
3. Remove the drive using one of the following methods:
 - To remove individual drives, click **Actions** for the drive you want to remove.
 - To remove multiple drives, select all the drives you want to remove and click **Bulk Actions**.
4. 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 Element OS web 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 136

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.

Replacing failed drives in the chassis

After you remove a failed drive from the failed drive list in the Element OS web UI, you are ready to physically replace the failed drive from the chassis.

Steps

1. Press on the latch mechanism and carefully pull the drive carrier from the slot.
2. Remove the failed drive.
3. Place the replacement drive in the empty slot.
4. Press firmly to seat it into the backplane.
5. Close the latch mechanism.

After the drive is seated, the drive should show up in the Element OS web UI list of available drives on Drives page.

6. Use the API or web UI to add it back to the cluster for it to become active.

Troubleshooting nodes

You can remove nodes from a cluster for maintenance or replacement. You should use the Element OS web 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 node.
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.

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.
- Shut down all the nodes at the same time.

Steps

1. Navigate to the management virtual IP (MVIP) address on the cluster to open the web 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

You can use the per-node utilities for troubleshooting network problems if the standard monitoring tools in the Element OS web UI do not give you enough information for troubleshooting. Per-node utilities provide specific information that can help you troubleshoot network problems between nodes or with the management node.

Related tasks

[Accessing per-node settings](#) on page 147

[Running system tests](#) on page 140

[Running system utilities from the per-node UI](#) on page 141

Related references

[Per-node network settings details](#) on page 139

[Per-node cluster settings details](#) on page 139

Per-node network settings details

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

The network settings for a node appear in the Network Settings window. You can see the network settings when you log in to the node. The following list describes the settings that you can modify when a node is in available, pending, and active states:

Method

The method used to configure the interface. This depends on other settings, such as the use of a static IP address, which will change the method to static. Possible methods are:

- `loopback`: Used to define the IPv4 loopback interface.
- `manual`: Used to define interfaces for which no configuration is done by default.
- `dhcp`: Might be used to obtain an IP address via Dynamic Host Configuration Protocol (DHCP).
- `static`: Used to define Ethernet interfaces with statically allocated IPv4 addresses.

IP Address

IP address for the 10G or 1G network.

Subnet Mask

Address subdivisions of the IP network.

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 bytes.

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`

Per-node cluster settings details

You can modify the cluster settings for a node. You can view the Cluster Settings tab after you log in to the node.

The following list describes the cluster settings for a node that you can modify:

Role

Role the node has in the cluster; it can be one of the following 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; it can be one of the following 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 OS running on the node.

Ensemble

Node 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.

Running system tests

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 node is stable and can be brought online without any issues.

Before you begin

You have logged in to the per-node user interface (UI).

Steps

1. Click **System Tests**.
2. Click the button to run the type of test that you want.
 - **Run All Tests**: Runs all test operations. This 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.
 - **Test Connect Mvip**: Pings the specified management virtual IP (MVIP) address and then executes a simple API call to the MVIP to verify connectivity.

- **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.
- **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. This test is resource intensive and should only be run if requested by 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:** Locates the cluster on the network by its name.
- **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 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 remotely paired clusters' nodes 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 from the per-node UI

You can use the per-node user interface (UI) 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 a per-node user interface (UI).

Steps

1. Click **System Utilities**.
2. Click the button to run the type of utility that you want.
 - **Create Support Bundle:** Creates a support bundle under the node's `/tmp/bundles` directory.
 - **Delete All Support Bundles:** Deletes any current support bundles residing on the node.
 - **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.
 - **Restart Networking:** Restarts all networking services on a node.

Attention: This operation can cause temporary loss of network connectivity.
 - **Restart Services:** Restarts Element OS 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.

Working with the management node

You can use the management node UI to enable remote access to your SolidFire cluster.

Related tasks

[Accessing a management node](#) on page 142

[Testing the management node settings](#) on page 143

[Running system utilities from the management node](#) on page 144

[Configuring a proxy server for the management node](#) on page 144

[Enabling remote NetApp Support connections](#) on page 145

Related references

[Management node network settings](#) on page 142

Accessing a management node

You can access network and cluster settings, and system tests and utilities in the management node user interface after you enter the management node IP and authenticate.

About this task

You need only one management node for reporting to and managing upgrades for a cluster. However, you might need to have multiple management nodes to allow multiple vCenter plug-ins to connect to a single cluster.

Steps

1. In a browser window, enter the management node IP address followed by 442:

```
https://192.168.107.112:442
```

2. In the authentication dialog box, enter an admin user name and password, if required.

After you log in, you can configure the management node and modify management node settings.

Related references

[Management node network settings](#) on page 142

Management node network settings

On the Network Settings tab of the Node UI for management node, you can modify the management node network interface fields.

Method

The method used to configure the interface. Possible methods are:

- `loopback`: Used to define the IPv4 loopback interface.
- `manual`: Used to define interfaces for which no configuration is done by default.
- `dhcp`: May be used to obtain an IP address via DHCP.
- `static`: Used to define Ethernet interfaces with statically allocated IPv4 addresses.

IP Address

IP address for the eth0 network.

Subnet Mask

Address subdivisions of the IP network.

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.

DNS Servers

Network interface used for cluster communication.

Search Domains

Search for additional MAC addresses available to the system.

Status

Possible values:

- UpAndRunning
- Down
- Up

Routes

Static routes to specific hosts or networks via the associated interface the routes are configured to use.

Management node cluster settings

On the Cluster Settings tab of the Node UI for the management node, you can modify cluster interface fields when a node is in Available, Pending, PendingActive, and Active states.

Role

Role the management node has in the cluster. Possible value: Management.

Hostname

Name of the management node.

Version

Element OS version running on the cluster.

Default Interface

Default network interface used for management node communication with the SolidFire cluster.

Testing the management node settings

After you change settings for the management node and commit the changes, you can run tests to validate the changes you made.

Steps

1. In the management node user interface, click **System Tests**.
2. To run all test operations, click **Run All Tests**.

3. To verify that the network settings you configured match the network settings being used on the system, click **Test Network Config**.
4. To check connectivity with host, click **Test Ping**, and specify the list of hosts to ping.
If you do not specify a host, the system pings each registered node in the cluster.

Related references

[Management node network settings](#) on page 142

Running system utilities from the management node

You can use the management node user interface (UI) to create or delete cluster support bundles, reset node configuration settings, or restart networking.

Before you begin

You are logged in to the management node using the management node admin credentials.

Steps

1. In the management node UI, click **System Utilities**.
2. Click the button to run the type of utility that you want.
 - **Create Cluster Support Bundle:** Creates the cluster support bundle to assist NetApp Support diagnostic evaluations of one or more nodes in a cluster. Click **show options** to display the fields you must fill.
 - **Delete All Support Bundles:** Deletes any current support bundles residing on the node.
 - **Reset Node:** Resets the management node to a new install image. This changes all settings to the default state, except the network configuration.
 - **Restart Networking:** Restarts all networking services on the management node.

Attention: This operation causes temporary loss of networking connectivity.

Configuring a proxy server for the management node

If your SolidFire cluster is behind a proxy server, you must configure the proxy settings so you can reach a public network.

Steps

1. Run the `sfsetproxy` to configure proxy settings for the management node.
The `sfsetproxy` command modifies settings for the following:
 - Reverse Secure Shell (SSH) tunnel
 - apt-get and aptitude (via `/etc/profile.d/sf_proxy_settings.sh`)
 - apt-mirror (via `wget`)
 - collector (script)
2. To verify that the proxy settings are consistent, run `sfsetproxy` without any arguments.

Example

For example, run the following command:

```
admin@mnode:~$ sudo sfsetproxy
```

The above command gives the following output:

```
Proxy host: 192.168.140.136
Proxy port: 3128
```

3. To set the host and port arguments when a user name and password are not required on the proxy server, run the following command:

```
sfsetproxy [-P ssh_port]
```

Example

```
sfsetproxy 192.168.140.136 3128
```

4. To set the host and port arguments when a user name and password are required on the proxy server, run the following command:

```
sfsetproxy [-P ssh_port] [-u username -p password]  
ProxyHostnameOrIPAddress ProxyPort (Initial setup of proxy)
```

Example

```
sfsetproxy -u testproxy -p solidfire 192.168.140.136 3128
```

Note: This command does not return output if it completes successfully. Run `./sfsetproxy` to see if the proxy has been set.

Enabling remote NetApp Support connections

If you require technical support for your SolidFire storage system, NetApp Support can connect remotely with your system. To gain remote access, NetApp Support can open a reverse Secure Shell (SSH) connection to your environment.

About this task

You can open a Transmission Control Protocol (TCP) port for SSH reverse tunnel connection with NetApp Support. This connection allows Support to log in to your management node. If your management node is behind a proxy server, the following TCP ports are required in the `sshd.config` file:

TCP port	Description	Connection direction
443	API calls/HTTPS for reverse port forwarding via open support tunnel to the web UI	Management node to SolidFire nodes
22	SSH login access	Management node to SolidFire nodes or from SolidFire nodes to management node

Steps

1. Log in to your management node and open a terminal session.

2. At a prompt, enter the following:

```
rst -r sfsupport.solidfire.com -u element -p <port number>
```

NetApp Support can provide the port number to access your management node with an SSH connection.

3. To close a remote support tunnel, enter the following:

```
rst --killall
```

Related references

[Network port requirements](#) on page 11

Understanding cluster fullness levels

The SolidFire cluster 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 Element OS web 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 OS web 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 web 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 tab of the Reporting page in the Element OS web UI.

The following list includes information about the MetadataClusterFull severity levels:

Warning

When the cluster is in this state, if two nodes are lost, there will not be enough capacity in the cluster to rebuild Double Helix data protection.

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.

Accessing per-node settings

You can access network and 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 admin user.

Note: You should configure or modify one node at a time. You must ensure that the network settings specified are having the expected effect, are stable, and are performing well before you make modifications to another node.

Step

1. Open the per-node UI by 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 OS web 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.

Enabling FIPS 140-2 on your cluster

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

About this task

Starting with Element OS 10.3, 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 Element OS Web 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.

If needed, contact NetApp Support for assistance.

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 references

[SSL ciphers](#) on page 148

SSL ciphers

SSL ciphers are encryption algorithms used by hosts to establish a secure communication. There are standard ciphers that Element OS 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 OS 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 on your cluster](#) on page 147

Enabling and disabling support access

You can use the SolidFire API to 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.

For detailed information, see the *NetApp SolidFire Element OS API Reference Guide*.

Where to find additional information

You can learn more about tasks related to the SolidFire system in NetApp's extensive documentation library.

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

Describes how to calculate the max provisioned space in a SolidFire cluster. Max provisioned space is the metadata drive capacity in a cluster.

[How to calculate SolidFire system error alert percentage](#)

Describes how to calculate the SolidFire cluster fullness percentage at which the system notifies you.

[NetApp SolidFire Element OS Documentation Library](#)

Includes documentation for various releases of Element OS.

[NetApp SolidFire Element OS Known Issues](#)

Provides a list of known issues in the current and earlier versions of Element OS.

[NetApp SolidFire Element OS Resolved Issues](#)

Provides a list of resolved issues in the current and earlier versions of Element OS.

[NetApp SolidFire Resources page](#)

Provides resources about SolidFire, including links to video content and technical reports.

[NetApp Technical Report 4677: NetApp SolidFire Element OS Remote Replication](#)

Describes different types of remote replication supported by NetApp SolidFire clusters.

Copyright information

Copyright © 2018 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark information

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>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 495 East Java Drive, Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277

Index

A

- access groups
 - adding a single initiator [81](#)
 - adding multiple initiators [82](#)
 - adding volumes [79](#)
 - creating [77](#)
 - deleting [83](#)
 - removing initiators [82](#)
 - removing volumes [79](#)
 - viewing details [79](#)
- access groups and initiators
 - working with [77](#)
- accounts
 - changing CHAP [55](#)
 - creating [54](#)
 - deleting [56](#)
 - editing [55](#)
 - finding individual account data [55](#)
 - overview [54](#)
 - understanding system data [55](#)
- administrator user accounts [27](#)
- alert error codes [129](#)
- alerts
 - error codes [129](#)
 - viewing [128](#), [131](#)
- API log
 - viewing API calls [24](#)

B

- blockServiceUnhealthy [135](#)
- broadcast client
 - enabling [35](#)

C

- cluster
 - mixed nodes [39](#)
- cluster admin accounts
 - changing password [30](#)
 - creating [28](#)
 - deleting [30](#)
 - LDAP [28](#)
- cluster admin permissions
 - editing [29](#)
- cluster administrator users
 - cluster admin [27](#)
 - creating [27](#)
 - deleting [27](#)
 - editing [27](#)
 - managing [27](#)
 - primary cluster admin [27](#)
- cluster admins [28](#)
- cluster capacity [39](#)
- cluster data management
 - overview of functions [54](#), [56](#)
- cluster fullness

- setting [32](#)
- cluster pairs
 - about cluster pairs [97](#)
 - deleting [100](#)
 - pairing with multiple targets [97](#)
 - real-time replication in local and remote pair [97](#)
- cluster settings
 - configuring [32](#)
 - configuring cluster fullness threshold
 - enabling support access [32](#)
 - encryption at rest [32](#)
 - NTP broadcast client [32](#)
 - SnapMirror [32](#)
 - virtual volumes [32](#)
- cluster upgrade
 - readiness checks for [48](#)
- clusters
 - about [9](#)
 - adding drives to [21](#)
 - creating new [20](#)
 - data refresh interval [25](#)
 - powering down [138](#)
 - QoS performance curve [58](#)
 - switch configuration for [10](#)
- comments
 - how to send feedback about documentation [153](#)
- connection.json
 - configuring [15](#)
 - setting up [15](#)
 - viewing help options [16](#)

D

- data collector
 - managing [16](#)
- data protection [84](#)
- Dell iDRAC Enterprise [19](#)
- details [113](#)
- disaster recovery
 - performing a failover [116](#)
 - SnapMirror [116](#)
- documentation
 - how to receive automatic notification of changes to [153](#)
 - how to send feedback about [153](#)
- drives
 - managing [37](#)
 - removing [37](#)
 - removing failed [135](#)
 - removing from chassis [137](#)
 - replacing [137](#)
 - resetting [141](#)
 - wear remaining [38](#)

E

- Element OS services
 - restarting [141](#)

Element OS software
 updating [50, 52](#)
 upgrading [45](#)
 upgrading on dark sites [51](#)
 upgrading using management node [45](#)

Element OS web UI
 accessing [23](#)
 accessing using MVIP [23](#)
 data refresh interval [25](#)
 filtering information [23](#)
 icons [25](#)
 providing feedback [26](#)
 sorting information [23, 24](#)
 sorting lists [24](#)
 sorting on multiple columns [24](#)
 sorting on single column [24](#)
 understanding cluster fullness [146](#)
 using [23](#)
 using sort criteria [24](#)
 viewing API calls [24](#)
 viewing the API log [24](#)

enabling SnapMirror on the volume [109](#)

encryption
 disabling encryption at rest [33](#)
 enabling encryption at rest [33](#)

encryption at rest
 disabling [33](#)
 enabling [33](#)

endpoints [109](#)

ESXi host details [76](#)

event log
 event types [126](#)
 using [125](#)
 viewing [125](#)

event types [126](#)

F

failback scenarios [119](#)

FC nodes
 configuring [19](#)

FC sessions
 information about [134](#)
 viewing [133](#)

feedback
 how to send comments about documentation [153](#)

Fibre Channel nodes
 configuring [19](#)

Fibre Channel ports
 viewing details [41](#)

Fibre Channel sessions
 information about [134](#)
 viewing [133](#)

G

group snapshots
 creating [90](#)
 deleting [91](#)
 editing [91](#)
 editing members of [92](#)

H

HealthTools
 checking installed version [46](#)
 managing on dark sites [46](#)
 performing software upgrades using [45](#)
 updating [47](#)

I

information
 how to send feedback about improving documentation [153](#)

initiators
 adding to a volume access group [81, 82](#)
 changing the alias [81](#)
 creating [80](#)
 deleting [83](#)
 editing [81](#)
 removing from access groups [82](#)

iSCSI sessions
 viewing [132](#)

L

LDAP
 configuring [31](#)
 configuring at the cluster level [30](#)
 creating cluster admin account [28](#)
 disabling [32](#)
 enabling authentication [31](#)
 managing [30](#)
 viewing settings [30](#)

LDAP settings
 configuring [27](#)

LDAP users and groups
 authorizing [30](#)

login banner [34](#)

M

management information base files
 downloading [37](#)
 MIB [37](#)
 viewing [37](#)

management node
 accessing [142](#)
 cluster interface settings [143](#)
 configuring [142](#)
 configuring proxy server [144](#)
 creating cluster support bundle [144](#)
 deleting cluster support bundle [144](#)
 for upgrading Element OS software [45](#)
 image types [14](#)
 installing [14](#)
 resetting configuration settings [144](#)
 restarting networking [144](#)
 upgrading software [49](#)

management node settings
 testing [143](#)

managing cluster administrator users [27](#)

managing cluster settings [27](#)

MDSStroubleshooting [135](#)**MDSS drives**adding [136](#)removing [37](#), [137](#)**MIB files**viewing [35](#)**mixed nodes**cluster capacity [39](#)**monitoring**using per-node utilities [138](#)**monitoring system** [125](#)**multi-drive slice service** [136](#), [137](#)**multiple volumes**cloning [93](#)cloning from a group snapshot [93](#)rolling back to a group snapshot [92](#)**N****NetApp Support**enabling remote connections [145](#)**network port requirements** [11](#)**network settings**testing changes [140](#)**network settings for eth0** [142](#)**network time protocol** [35](#)**networking**overview [10](#)**nodes**accessing settings using per-node UI [147](#)adding to a cluster [38](#)changing cluster settings [139](#)changing network settings [139](#)configuring [17](#)configuring iDRAC [19](#)configuring using Element OS [17](#)configuring using node UI [18](#)configuring using TUI [17](#), [19](#)Element OS software version [39](#)Fibre Channel [10](#)managing [38](#)managing hardware [19](#)monitoring hardware [19](#)overview [9](#)removing [137](#)resetting [141](#)resetting configuration settings [144](#)running system tests [140](#)running system utilities [141](#)states [40](#)storage [9](#)troubleshooting [137](#)upgrading Element OS software [45](#)version numbers [39](#)versioning [39](#)viewing details for individual [41](#)viewing performance activity [130](#)**NTP** [35](#)**P****paired clusters**deleting [100](#)**paired volumes**deleting [107](#)editing [107](#)**pairing clusters**pairing using MVIP [98](#)pairing using pairing key [98](#)pairing without remote cluster access [98](#)**passwords**cluster admin account [30](#)**performing a failback**SnapMirror [117](#)**protocol endpoints information** [75](#)**Q****QoS policies**about [59](#)applying to volumes [67](#)changing QoS policy settings [60](#)creating [59](#)deleting a QoS Policy [60](#)disassociating a policy from a volume [67](#)viewing details [60](#)**Quality of Service (QoS)**about burst IOPS [57](#)about maximum IOPS [57](#)about minimum IOPS [57](#)

maximum possible values

minimum possible values [58](#)performance curve [58](#)**R****relationships**actions [115](#)**replication**assigning a replication source volume [103](#)assigning a replication target [103](#)cluster pair details [99](#)cluster pairing [96](#), [97](#)cluster pairs [97](#)messages during volume pairing [105](#)pairing volumes [100](#)pairing volumes using a volume ID [101](#), [102](#)real-time replication [96](#)viewing cluster pair details [99](#)viewing volume pair details [104](#)volume pair [100](#)volume pair details [104](#)volume pairing [96](#), [100](#)volume pairs [100](#)warnings during volume pairing [106](#)**running tasks**viewing progress of [128](#)viewing status of [128](#)

S

SnapMirror
 adding labels to snapshots [111](#)
 creating endpoints [109](#)
 deleting endpoints [111](#)
 editing an endpoint [110](#)
 enabling on the cluster [108](#)
 enabling on the volume [109](#)
 endpoint details [110](#)
 endpoints [109](#)
 failback [119](#)
 labels [111](#)
 manually enabling [108](#)
 performing a failback [117](#), [118](#)
 performing a failover [116](#)
 retention rules [111](#)
 source volume still exists [117](#)
SnapMirror failback [117](#)
SnapMirror labels
 adding to snapshot schedules [112](#)
 adding to snapshots [111](#)
 specifying retention policies [111](#)
SnapMirror relationships
 creating [112](#)
 deleting [115](#)
 editing [114](#)
 managing [108](#)
 modifying [114](#)
 snapshot schedules
 copying [96](#)
 creating [94](#)
 deleting [96](#)
 editing [95](#)
SNMP
 configuring [35](#)
 configuring requestors [35](#)
 configuring traps [36](#)
 configuring USM users [36](#)
 managing [35](#)
 overview [35](#)
 solid-state disk drives
 troubleshooting [134](#)
SolidFire Active IQ
 connecting [15](#)
SolidFire and ONTAP
 replicating between clusters [108](#)
SolidFire system
 additional documentation about [150](#)
 management [27](#)
 overview [9](#)
 setting up [14](#)
 source volume
 resynchronizing [117](#), [118](#)
 source volume lost
 performing a failback [118](#)
 source volume still exists
 performing a failback [117](#)
SSDs
 troubleshooting [134](#)
SSL certificate [22](#)
SSL ciphers [148](#)
 storage containers

deleting [74](#)
 storage nodes
 configuring [17](#)
 configuring using Element OS [17](#)
 configuring using TUI [17](#)
 suggestions
 how to send feedback about documentation [153](#)
 support access
 disabling [149](#)
 enabling [149](#)
 support bundle
 creating [141](#)
 system tests [140](#)

T

terminal user interface [19](#)
Terms of Use
 configuring [34](#)
 configuring banner [34](#)
 disabling [34](#)
 editing [34](#)
 enabling [34](#)
 managing [34](#)
 troubleshooting
 using per-node utilities [138](#)
 using the management node UI [142](#)
 troubleshooting system [125](#)
TUI [19](#)
Twitter
 how to receive automatic notification of
 documentation changes [153](#)

U

upgrading software [27](#)
 user guide
 overview [8](#)
 user types [27](#)

V

virtual networks
 adding [43](#)
 deleting [45](#)
 editing [44](#)
 managing [42](#)
 viewing details [42](#)
 virtual routing and forwarding
 enabling [44](#)
 VLANs [44](#)
 virtual volumes
 about bindings [76](#)
 about storage containers [73](#), [75](#)
 bindings information for each virtual volume [76](#)
 changing CHAP authentication on a storage
 container [74](#)
 configuring VVols [68](#)
 creating a storage container [73](#)
 enabling functionality on the cluster [68](#)
 individual storage container details [74](#)
 storage container details [73](#)

- using the Element OS to delete a VVol [72](#)
- using the UI to manage [68](#)
- viewing high-level VVol performance activity [70](#)
- viewing individual VVol performance activity [71](#)
- viewing VVol performance activity [69](#)
- volume access groups [77, 78](#)
- volume backups
 - restoring from a SolidFire cluster [124](#)
 - restoring from Amazon S3 [122](#)
 - restoring from OpenStack Swift [123](#)
- volume pairs
 - assigning a replication source [103](#)
 - assigning a replication target [103](#)
 - deleting [107](#)
 - editing [107](#)
 - messages during volume pairing [105](#)
 - warnings during volume pairing [106](#)
- volume performance [131](#)
- volume snapshot backup operations [87](#)
- volume snapshots
 - backing up to a SolidFire cluster [88](#)
 - backing up to an Amazon S3 object store [87](#)
 - backing up to an OpenStack Swift object store [88](#)
 - cloning volumes from [86](#)
 - creating [84](#)
 - deleting [86](#)
 - editing the retention period for [85](#)
 - rolling back individual volumes to [87](#)
- volumes
 - adding to access groups [79](#)
 - assigning a replication source [103](#)
 - assigning a replication target [103](#)
 - assigning LUNs to Fibre Channel volumes [67](#)
 - assigning QoS during volume creation [61](#)
 - backing up and restoring [120](#)
 - backing up to a SolidFire cluster [121](#)
 - backing up to Amazon S3 [120](#)
 - backing up to OpenStack Swift [121](#)
 - changing associated account for a volume [63](#)
 - changing volume access [63](#)
 - changing volume access mode [63](#)
 - changing volume QoS [63](#)
 - changing volume size [63](#)
 - cloning [65](#)
 - creating a volume [61](#)
 - deleting [64](#)
 - deleting a volume with an associated snapshot [64](#)
 - details [61](#)
 - editing volume settings [63](#)
 - length of time before purge [64](#)
 - pairing using a pairing key [102](#)
 - pairing using a volume ID [101](#)
 - purging [65](#)
 - removing from access groups [79](#)
 - restoring a deleted volume [65](#)
 - viewing individual volume performance [62](#)
- VRF [44](#)
- VRF VLANs
 - editing [44](#)