

NetApp Element 11.5

Element Plug-in for vCenter Server User Guide

Version 4.3

November 2019 | 215-14588_2019-11_en-us
doccomments@netapp.com

NetApp®

Contents

About this guide 8
vCenter Plug-in overview 9
Network port requirements 10
VMware vSphere prerequisites 14
Getting started 15
Installing the NetApp Element Plug-in for vCenter Server 15
Upgrading the vCenter Plug-in 15
Upgrading vCenter Server 18
Installing a management node 18
Registering the vCenter Plug-in with vCenter 22
Modifying vCenter properties for an in-house (dark site) HTTP server 25
Upgrading the management node 28
Migrating from management node 10.x to version 11.x 28
Upgrading the management node to version 11.5 from version 11.3 30
Upgrading the management node to version 11.5 from version 11.0 or 11.1 32
Configuring a storage NIC (eth1) 35
Persistent volumes 36
Updating the vCenter Plug-in registration for vCSA 36
Updating the vCenter plug-in registration for Windows vCenter 39
Updating management services 41
Updating management services with NetApp Hybrid Cloud Control 42
Updating management services using mNode API 42
Accessing the plug-in after installation 44
How to use the NetApp Element Plug-in for vCenter Server 45
NetApp Element Configuration Extension Point 45
NetApp Element Management Extension Point 46
vCenter Linked Mode 48
Using object naming best practices when managing multiple clusters 49
Cluster configuration 50
Adding a cluster 50
Viewing cluster details 51
Cluster details 51
Individual cluster details 51
Editing cluster profiles 52
Removing a cluster profile 53
Enabling virtual volumes 53
Enabling encryption at rest 54
Disabling encryption at rest 54
Enabling SSH 55

Disabling SSH 55
Changing the SSH time limit 55
Setting protection domain monitoring 56
Shutting down a cluster 56
Expanding your NetApp HCI 57
Management node (mNode) settings 58
Configuring management node settings for QoSSIOC 58
Editing management node settings for QoSSIOC 59
Clearing management node settings for QoSSIOC 59
Changing the QoSSIOC service password 60
Viewing QoSSIOC events 62
QoSSIOC event details 62
Plug-in product information 64
Reporting 65
Reporting overview 65
Viewing event logs 67
Event log 67
Event types 68
Alerts 69
Alert error codes 70
iSCSI sessions 76
Running tasks 77
Volume performance 78
Management 80
Datastore management 80
Creating a datastore 80
Viewing the datastore list 83
Extending a datastore 84
Cloning a datastore 84
Sharing a datastore 87
Performing VAAI UNMAP 88
Deleting a datastore 89
QoSSIOC automation 89
Enabling QoSSIOC automation 90
Disabling QoSSIOC integration 91
Volume management 91
Creating a volume 91
Viewing volumes details 93
Editing a volume 97
Cloning a volume 98
Volume backup and restore operations 99
Deleting volumes 105
Purging volumes 106
Restoring deleted volumes 106
Adding volumes to an access group 107

Removing volumes from an access group 107
Applying a QoS policy to multiple volumes 108
Changing the QoS policy association of a volume 108
User account management 109
Creating an account 109
Account details 109
Editing an account 110
Deleting an account 110
Volume access groups 111
Creating access groups 111
Volume access group details 112
Editing access groups 112
Deleting access groups 113
Initiators 113
Creating an initiator 113
Initiator details 114
Editing an initiator 114
Deleting initiators 115
Adding initiators to a volume access group 115
QoS policies 115
Creating a QoS policy 116
QoS policy details 116
Editing a QoS policy 117
Deleting a QoS policy 117
Data protection 118
Volume snapshots 118
Creating a volume snapshot 118
Volume snapshot details 119
Editing snapshots 120
Cloning a volume from a snapshot 121
Rolling back a volume to a snapshot 122
Volume snapshot backup operations 122
Deleting a snapshot 125
Group snapshots 126
Creating a group snapshot 126
Group snapshot details 127
Editing group snapshots 127
Cloning volumes from a group snapshot 128
Rolling back volumes to a group snapshot 128
Deleting a group snapshot 129
Snapshot schedules 130
Creating a snapshot schedule 130
Snapshot schedule details 131
Editing a snapshot schedule 131
Copying a snapshot schedule 132

Deleting a snapshot schedule 133
Configuring cluster and volume pairing for real-time remote replication 133
Pairing clusters 134
Pairing volumes 137
Validating volume replication 145
Deleting a volume relationship after replication 145
Managing volume relationships 145
Cluster management 151
Drives 151
Adding available drives to a cluster 151
Drive details 152
Removing a drive 153
Nodes 153
NetApp HCI storage nodes 153
NetApp HCI compute nodes 153
SolidFire Storage nodes 153
SolidFire Fibre Channel nodes 154
Adding a node to a cluster 154
Node details 155
Removing nodes from a cluster 156
Restarting cluster nodes 157
Shutting down cluster nodes 157
VLAN management 158
Creating a VLAN 158
Virtual network details 159
Editing a virtual network 160
Deleting a virtual network 160
Virtual volumes 161
Configuring vSphere Virtual Volumes (VVols) functionality 161
Registering the NetApp Element VASA Provider 161
Creating a VVol datastore 162
Viewing virtual volume details 163
Virtual volume details 164
Individual virtual volume details 164
Storage containers 168
Creating a storage container 168
Viewing storage container details 169
Deleting a storage container 172
Protocol endpoints 172
Viewing protocol endpoint details 172
Unregistering the vCenter Plug-in 175
Removing the vCenter Plug-in 177
Troubleshooting 179
Plug-in registration successful but icons do not appear in web client 179
Error registering plug-in using Registration UI 180

Error updating plug-in using Registration UI 180
Keystore error using Registration UI 180
Error message that NetApp extension cannot be upgraded 181
Removing plug-in completes successfully but icons remain 181
Plug-in cannot be unregistered or removed after admin password change 181
Plug-in management tasks fail or volumes are not accessible to ESXi host 182
Failure occurs during vCenter Plug-in use on Firefox 59.0.2 browsers 182
Delete datastore operation fails 182
Cluster pair cannot connect using a pairing key 183
Error message for QoSSIOC status 183
QoSSIOC service shown as available but is unavailable 183
QoSSIOC is enabled for datastore but unavailable 184

Where to find product documentation and other information 185
Contacting NetApp Support 186
Copyright 187
Trademark information 188
How to send comments about documentation and receive update notifications 189

About this guide

This guide provides information about how to configure and manage storage clusters using the NetApp Element Plug-in for vCenter Server. The intended audience for this guide is those who install, administer, or troubleshoot storage solutions, and VMware admins who need to allocate storage for virtual machines (VMs). Other IT professionals or software developers may also find this document useful.

vCenter Plug-in overview

The NetApp Element Plug-in for vCenter Server (VCP) is a web-based tool integrated with the VMware vSphere Web Client user interface (UI). The plug-in is an extension and alternative scalable, user-friendly interface for VMware vSphere that can manage and monitor storage clusters running NetApp Element software.

You can use the plug-in user interface to discover and configure clusters, and to manage, monitor, and allocate storage from cluster capacity to configure datastores and virtual datastores (for virtual volumes). A cluster appears on the network as a single local group that is represented to hosts and administrators by virtual IP addresses. You can also monitor cluster activity with real-time reporting, including error and alert messaging for any event that might occur while performing various operations.

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 Support Server, management node, and the storage nodes.

Tip: Enable ICMP between the management node, storage nodes, and cluster MVIP.

Note: For vSphere network port requirements, see VMware documentation.

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 and Element software upgrades
Management node	Online software repository: • `https://repo.netapp.com/bintray/api/package` • `https://netapp-downloads.bintray.com`	443	Management node service upgrades
Management node	`monitoring.solidfire.com`	443	Storage cluster reporting to Active IQ
Management node	Storage cluster MVIP	443	UI and API access to storage node and Element software upgrades

Source	Destination	Port	Description
Management node	`connect.pub.nks.cloud`	443	Provides secure communications between NKS Cloud Providers and the hosted NKS Service, for example, when NKS is deployed on NetApp HCI or VMware on-premises traffic makes use of this Northbound MTLS secured channel.
Management node	`api.nks.netapp.io`	443	Facilitates initial deployment-time registration of on-premises "regions."
Management node	`repo.netapp.com`	443	Provides access to components necessary to install/update on-premises deployment.
34.208.181.140 34.217.162.31 54.187.65.159 18.236.231.155	Management node	443	HTTPS (Kubernetes cluster security).
		6443	Kubernetes API (Kubernetes cluster security).
		12443	Proxy to dashboard (Kubernetes cluster security).
		22	Kubernetes upgrades and other tasks (Kubernetes cluster security).
Management node	`amazonaws.com`	443	Dispatch tunnel
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 software upgrades
Storage node MIP	S3/Swift endpoint	80	(Optional) HTTP communication to S3/Swift endpoint for backup and recovery
Storage node MIP	NTP server	123 UDP	NTP
Storage node MIP	Management node	162 UDP	(Optional) SNMP traps
Storage node MIP	SNMP server	162 UDP	(Optional) SNMP traps
Storage node MIP	LDAP server	389 TCP/UDP	(Optional) 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

Source	Destination	Port	Description
Storage node MIP	S3/Swift endpoint	443	(Optional) HTTPS communication to S3/Swift endpoint for backup and recovery
Storage node MIP	Management node	10514 TCP/UDP 514 TCP/UDP	Syslog forwarding
Storage node MIP	Syslog server	10514 TCP/UDP 514 TCP/UDP	Syslog forwarding
Storage node MIP	LDAPS server	636 TCP/UDP	LDAPS lookup
Storage node MIP	Remote storage node MIP	2181	Intercluster communication for remote replication
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	Intercluster communication for remote replication
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
Storage node SIP	Compute node SIP	442	Compute node API, configuration and validation, and access to software inventory
System administrator PC	Storage node MIP	80	(NetApp HCI only) Landing page of NetApp Deployment Engine
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
			(NetApp HCI only) Configuration and deployment monitoring in NetApp Deployment Engine

Source	Destination	Port	Description
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
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	8443	(Optional) vCenter Plug-in QoSSIOC service.
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.

VMware vSphere prerequisites

vSphere 6.0, 6.5, or 6.7, including vCenter and ESXi with software iSCSI adapter and iSCSI networking configured, is required to use the NetApp Element Plug-in for vCenter Server.

The plug-in is not compatible with version 6.7 U2 of the HTML5 vSphere Web Client. It is compatible with the version 6.7 U2 vSphere Web Client for Flash/FLEX.

Because datastores are created using the highest VMFS version supported by the selected ESXi host, all cluster members should run the same version of vSphere and ESXi to avoid VMFS compatibility issues.

Attention: The vSphere HTML5 web client and Flash web client have separate databases that cannot be combined. Clusters added in one client will not be visible in the other. If you intend to use both clients, add your clusters in both.

Getting started

You can install the most recent version of the NetApp Element Plug-in for vCenter Server (VCP) directly to your vCenter and access the plugin with the vSphere Web Client. After installation is complete, you can utilize the quality of service based on storage I/O control (QoSSIOC) service as well as other services of the vCenter Plug-in.

To get started or update an existing installation, see one of the following topics:

- New installations: *Installing the NetApp Element Plug-in for vCenter Server*
- Existing installations: *Upgrading the vCenter Plug-in*

Installing the NetApp Element Plug-in for vCenter Server

For new installations, you can manually install the NetApp Element plug-in for vCenter Server by deploying a new management node and then registering the plug-in with a vCenter Server.

About this task

vCenter high availability (VCHA) is not supported.

Steps

1. *Download and deploy the OVA for your NetApp HCI or NetApp Element software installation using the vSphere Web Client*.
2. *Register the plug-in with vCenter*.
3. *Add storage clusters using the plug-in*.
4. *Configure management node and QoSSIOC settings using the plug-in*.
5. *Updating management services with NetApp Hybrid Cloud Control*

Related concepts

Accessing the plug-in after installation on page 44

Upgrading the vCenter Plug-in

For existing installations, you can upgrade the NetApp Element Plug-in for vCenter Server (VCP) using the latest upgrade package or by taking steps to migrate settings from the Flash to the HTML5 version of vSphere Web Client.

Before you begin

You have a cluster running NetApp Element software 11.3 or later.

Step

1. Upgrade your plug-in according to one of the following procedures if the statements described in the prerequisites are true about your installation and upgrade plans:

   Attention: Perform any required vCenter upgrades before upgrading the plug-in. These procedures assume that vCenter upgrades have already been completed.

Prerequisites	Steps
• Your current plug-in 3.0.1 or 4.*x* is registered with vCenter Server. • You are using the Flash-based vSphere Web Client. • You intend to upgrade your plug-in for use with a Flash version of vSphere Web Client.	**a.** Log out of the vSphere Web Client. **b.** Choose one of the following management node upgrade options: • If you are upgrading from management node 11.0, 11.1, or 11.3, download and deploy the management node ISO for Element software 11.5 or NetApp HCI 1.7 and complete one of the following the upgrade procedures: ◦ *Upgrading the management node to version 11.5 from version 11.3* ◦ *Upgrading the management node to version 11.5 from version 11.0 or 11.1* • If you are upgrading from a management node version 10.x, you must migrate your management node settings to a new management node 11.1 VM before you can update to management node 11.3 and later. vCenter Plug-in 4.3 requires management node 11.3 or later. *Migrating from management node 10.x to version 11.x* **c.** Using the registration utility (`https://[management node IP]:9443`), update the plug-in registration with vCenter from *vCSA* or *Windows*. **Note:** Modify properties for an in-house (dark site) HTTP server, if required. *Modifying properties for HTTP server* **d.** Update management services: *Updating management services with NetApp Hybrid Cloud Control* **e.** Log in to vSphere Web Client and confirm the upgrade was successful. *Accessing the plug-in after installation*

Prerequisites	Steps
• Your current plug-in 3.0.1 or 4.*x* is registered with vCenter Server. • You are using the Flash-based vSphere Web Client. • You intend to upgrade your plug-in for use with a version 6.5 or 6.7 HTML5 vSphere Web Client.	**a.** Log out of the vSphere Web Client. **b.** Choose one of the following management node upgrade options: • If you are upgrading from management node 11.0, 11.1, or 11.3, download and deploy the management node ISO for Element software 11.5 or NetApp HCI 1.7 and complete one of the following the upgrade procedures: ◦ *Upgrading the management node to version 11.5 from version 11.3* ◦ *Upgrading the management node to version 11.5 from version 11.0 or 11.1* • If you are upgrading from a management node version 10.x, you must migrate your management node settings to a new management node 11.1 VM before you can update to management node 11.3 and later. vCenter Plug-in 4.3 requires management node 11.3 or later. *Migrating from management node 10.x to version 11.x* **c.** Using the registration utility (`https://[management node IP]:9443`), update the plug-in registration with vCenter from *vCSA* or *Windows*. **Note:** Modify properties for an in-house (dark site) HTTP server, if required. *Modifying properties for HTTP server* **d.** Update management services: *Updating management services with NetApp Hybrid Cloud Control* **e.** Log in to vSphere HTML5 Web Client and confirm the upgrade was successful. *Accessing the plug-in after installation* **f.** Use the plug-in to complete the following steps: **i.** *Add storage clusters*. **ii.** *Configure management node and QoSSIOC settings*.
• Your current plug-in 4.1 or later is registered with vCenter Server. • You are using the Flash-based vSphere Web Client. • You intend to keep your existing plug-in version rather than upgrade. • You intend to switch from the Flash-based vSphere Web Client to the HTML5-based vSphere Web Client.	**a.** *Clear the management node settings in the vSphere Flash Web Client*. **b.** Log out of the vSphere Flash Web Client. **c.** Log in to the vSphere HTML5 Web Client. **d.** Use the plug-in to complete the following steps: **i.** *Add storage clusters*. **ii.** *Configure management node and QoSSIOC settings*.

Related concepts

Accessing the plug-in after installation on page 44

Related tasks

Unregistering the vCenter Plug-in on page 175
Removing the vCenter Plug-in on page 177
Registering the vCenter Plug-in with vCenter on page 22

Upgrading vCenter Server

You must unregister the vCenter Plug-in before making vCenter Server upgrades and register the plug-in again after the upgrade. Because vCenter Server upgrades dispose of plugin data, you must also reconfigure clusters and, if needed, management node (mNode) QoSSIOC settings.

Steps

1. Unregister the plug-in from the vCenter with which it is associated.
2. Upgrade vCenter to the latest version.
3. Register the plug-in with the upgraded vCenter.
4. Add clusters using the plug-in.
5. Optional: Configure management node settings.

Related tasks

Unregistering the vCenter Plug-in on page 175
Registering the vCenter Plug-in with vCenter on page 22
Adding a cluster on page 50
Configuring management node settings for QoSSIOC on page 58

Installing a management node

You can manually install the management node for your cluster running NetApp Element software using the appropriate image for your configuration. This manual process is intended for SolidFire all-flash storage administrators and NetApp HCI administrators who are not using the NetApp Deployment Engine for management node installation.

Before you begin

- Your cluster version must be running NetApp Element software 11.3 or later.
- Your installation uses IPv4. The management node 11.3 does not support IPv6.

  Note: You can use the management node 11.1 if you need IPv6 support.

- You have permissions to download software from the NetApp Support Site.
- 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`

Platform	Installation image type
KVM	`.iso`
VMware vSphere	`.iso`, `.ova`
Citrix XenServer	`.iso`
OpenStack	`.iso`

About this task

Prior to completing this procedure, you should have an understanding of persistent volumes and whether or not you want to use them. Persistent volumes allow management node data to be stored on a specified storage cluster so that data can be preserved in the event of management node loss or removal.

Steps

1. Download the OVA or ISO for your installation from the NetApp Support Site:
   - Element software: *https://mysupport.netapp.com/products/p/element_software.html*
   - NetApp HCI: *https://mysupport.netapp.com/products/p/hci.html*
   a. Select the version number of the software to download.
   b. Click **Go**.
   c. Read and click through the required prompts, accept the EULA, and select the management node image you want to download.
2. If you downloaded the OVA, follow these steps:
   a. Deploy the OVA.
   b. If your storage cluster is on a separate subnet from your management node (eth0) and you want to use persistent volumes, add a second network interface controller (NIC) to the VM on the storage subnet (eth1) or ensure that the management network can route to the storage network..
3. If you downloaded the ISO, follow these steps:
   a. Create a new 64-bit virtual machine from your hypervisor with the following configuration:
      - Six virtual CPUs
      - 12GB RAM
      - 400GB virtual disk, thin provisioned
      - One virtual network interface with internet access and access to the storage MVIP.
      - (Optional for SolidFire all-flash storage) One virtual network interface with management network access to the storage cluster. If your storage cluster is on a separate subnet from your management node (eth0) and you want to use persistent volumes, add a second network interface controller (NIC) to the VM on the storage subnet (eth1) or ensure that the management network can route to the storage network.

      Attention: Do not power on the virtual machine prior to the step indicating to do so later in this procedure.

   b. Attach the ISO to the virtual machine and boot to the `.iso` install image.

Note: Installing a management node using the image might result in 30-second delay before the splash screen appears.

4. Power on the virtual machine for the management node after the installation completes.
5. Using the terminal user interface (TUI), create a management node admin user.

   Tip: To enter text, press **Enter** three times 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.

6. Configure the management node network (eth0).

   Note: If you have a second NIC on eth1, see instructions on configuring a second NIC.

   Configuring a storage NIC (eth1)

7. SSH into the management node.
8. Using SSH, run the following command to gain root privileges. Enter your password when prompted:

   ```
   sudo su
   ```

9. Ensure time is synced (NTP) between the management node and the storage cluster.

   Note: In vSphere, the **Synchronize guest time with host** box should be checked in the VM options. Do not disable this option if you make future changes to the VM.

10. Configure the management node setup command:

    Note: You might be prompted to enter passwords or other information if you do not include them in the command. If your cluster is behind a proxy server, you must configure the proxy settings so you can reach a public network.

    ```
    /sf/packages/mnode/setup-mnode --mnode_admin_user [username] --
    storage_mvip [mvip] --storage_username [username] --telemetry_active
    [true]
    ```

    a. Replace the value in `[ ]` brackets (including the brackets) for each of the following required parameters:

       Note: The abbreviated form of the command name is in parentheses `( )` and can be substituted for the full name.

       --mnode_admin_user (-mu) [username]

       The username for the management node administrator account. This is likely to be the username for the user account you used to log into the management node.

       --storage_mvip (-sm) [MVIP address]

       The MVIP (management virtual IP address) of the storage cluster running Element software.

       --storage_username (-su) [username]

       The storage cluster administrator username for the cluster specified by the `--storage_mvip` parameter.

       --telemetry_active (-t) [true]

       Retain the value `true` that enables data collection for analytics by Active IQ.

b. (Optional): Add password or Active IQ endpoint parameters to the command. You will be prompted to enter these passwords in a secure prompt if you do not include them in the command:

--mnode_admin_password (-mp) [password]

The password of the management node administrator account. This is likely to be the password for the user account you used to log into the management node.

--storage_password (-sp) [password]

The password of the storage cluster administrator specified by the `--storage_username` parameter.

--remote_host (-rh) [AIQ_endpoint]

The endpoint where Active IQ telemetry data is sent to be processed. If the parameter is not included, the default endpoint is used.

c. (Optional): Add the following persistent volume parameters:

> **Attention:** Do not modify or delete the account and volumes created for persistent volumes functionality or a loss in management capability will result.

--use_persistent_volumes (-pv) [true/false, default: false]

Enable or disable persistent volumes. Enter the value `true` to enable persistent volumes functionality.

--persistent_volumes_account (-pva) [account_name]

If `--use_persistent_volumes` is set to `true`, use this parameter and enter the storage account name that will be used for persistent volumes.

> **Note:** Use a unique account name for persistent volumes that is different from any existing account name on the cluster. It is critically important to keep the account for persistent volumes separate from the rest of your environment.

--persistent_volumes_mvip (-pvm) [mvip]

Enter the MVIP (management virtual IP address) of the storage cluster running Element software that will be used with persistent volumes. This is only required if multiple storage clusters are managed by the management node. If multiple clusters are not managed, the default cluster MVIP will be used.

d. Configure a proxy server:

--use_proxy (-up) [true/false, default: false]

Enable or disable the use of the proxy. This parameter is required to configure a proxy server.

--proxy_hostname_or_ip (-pi) [host]

The proxy hostname or IP. This is required if you want to use a proxy. If you specify this, you will be prompted to input `--proxy_port`.

--proxy_username (-pu) [username]

The proxy username. This parameter is optional.

--proxy_password (-pp) [password]

The proxy password. This parameter is optional.

--proxy_port (-pq) [port, default: 0]

The proxy port. If you specify this, you will be prompted to input the proxy host name or IP (`--proxy_hostname_or_ip`).

--proxy_ssh_port (-ps) [port, default: 443]

The SSH proxy port. This defaults to port 443.

11. (Optional) Use parameter help if you need additional information about each parameter:

    --help (-h)

    Returns information about each parameter. Parameters are defined as required or optional based on initial deployment. Upgrade and redeployment parameter requirements might vary.

12. Run the `setup-mnode` command.

13. Use the mNode API to add assets:

    a. Using a browser, go to the storage MVIP and log in.

       This action causes certificate to be accepted for the next step.

    b. Using a browser, go to `https://<ManagementNodeIP>/mnode`.

    c. Add a vCenter controller asset to the management node known assets for HCI monitoring (NetApp HCI installations only) and Hybrid Cloud Control (HCC):

    d. Click **Authorize** and enter your MVIP user name and password credentials. Close the pop-up window.

    e. Run **GET /assets** to pull the base asset ID needed to add the vCenter/controller asset.

    f. Run **POST /assets/{ASSET_ID}/controllers** to add a controller asset with vCenter credentials.

    g. For NetApp HCI or to access cloud services options in HCC, add a compute asset to the management node known assets:

       Attention: You must perform this step or Cloud Services options will not be available from HCC.

    h. Click **Authorize** and enter your MVIP user name and password credentials. Close the pop-up window.

    i. Run **GET /assets** to pull the base asset ID needed to add the compute asset.

    j. Run **POST/assets/{asset_id}/compute-nodes** to add a compute asset with credentials for the compute asset. The type is `ESXi Host`.

Related tasks

Upgrading the vCenter Plug-in on page 15
Installing the NetApp Element Plug-in for vCenter Server on page 15

Registering the vCenter Plug-in with vCenter

You can deploy the vCenter Plug-in package in the vSphere Web Client by registering the package as an extension on vCenter Server. After registration is complete, the plug-in is available to any vSphere Web Client that connects to your vSphere environment.

Before you begin

- You have logged out of the vSphere Web Client.

  Note: The web client will not recognize updates made during this process if you do not log out.

- You have vCenter Administrator role privileges to register a plug-in.
- You have deployed a management node OVA running Element software 11.3 or later.

- Your management node is powered on with its IP address or DHCP address configured.
- You are using an SSH client or web browser (Chrome 56 or later or Firefox 52 or later).
- Your firewall rules allow open network communication between the vCenter and the storage cluster MVIP on TCP ports 443, 8443, and 9443. Port 9443 is used for registration and can be closed after registration is complete. If you have enabled virtual volumes functionality on the cluster, ensure TCP port 8444 is also open for VASA provider access.

About this task

You must register the vCenter Plug-in on every vCenter Server where you need to use the plug-in.

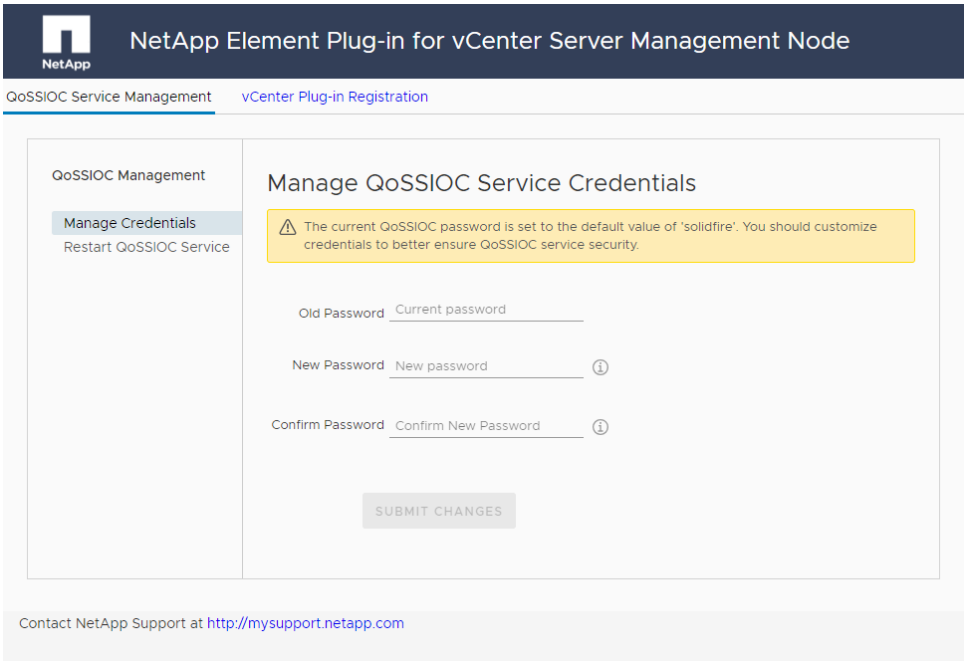
Note: The plug-in must be registered with each vCenter Server in a Linked Mode environment to keep MOB data in sync and to be able to upgrade the plug-in.

When a vSphere Web Client connects to a vCenter Server where your plug-in is not registered, the plug-in is not visible to the client.

Steps

1. Enter the IP address for your management node in a browser, including the TCP port for registration: `https://[management node IP]:9443`.

   The registration UI displays the **Manage QoSSIOC Service Credentials** page for the plug-in.



2. Optional: Change the password for the QoSSIOC service before registering the vCenter Plug-in:

   a. Enter the following information:

      - **Old Password**: The current password of the QoSSIOC service. If you have not yet assigned a password, type the default password:

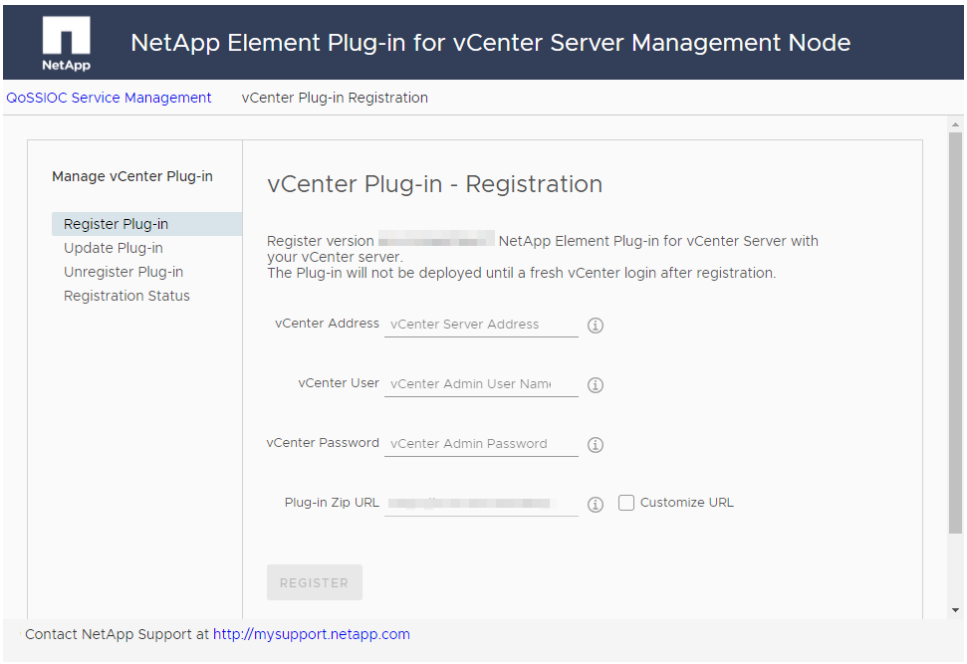
        `solidfire`

      - **New Password**: The new password for the QoSSIOC service.
      - **Confirm Password**: Enter the new password again.

b. Click **Submit Changes**.

Note: The QoSSIOC service automatically restarts after you submit changes.

3. Click **vCenter Plug-in Registration**.



4. Enter the following information:

   - The IPv4 address or the FQDN of the vCenter service on which you will register your plug-in.
   - The vCenter Administrator user name.

     Note: The user name and password credentials you enter must be for a user with vCenter Administrator role privileges.

   - The vCenter Administrator password.
   - (For in-house servers/dark sites) A custom URL for the plug-in ZIP.

     Note: Most installations use the default path. You can click **Custom URL** to customize the URL if you are using an HTTP or HTTPS server (dark site) or have modified the ZIP file name or network settings. For additional configuration steps if you intend to customize a URL, see plug-in documentation about modifying vCenter properties for an in-house (dark site) HTTP server.

5. Click **Register**.
6. Optional: Click **Registration Status**.
7. Optional: Enter the following information:

   - The IPv4 address or the FQDN of the vCenter service on which you are registering your plug-in
   - The vCenter Administrator user name
   - The vCenter Administrator password

8. Optional: Click **Check Status** to verify that the new version of the plug-in is registered on the vCenter Server.
9. Log in to the vSphere Web Client as a vCenter Administrator.

   Note: This action completes the installation in the vSphere Web Client. If the vCenter Plug-in icons are not visible from the vSphere main page, see documentation about troubleshooting the plug-in.

Related tasks

Upgrading the vCenter Plug-in on page 15
Installing the NetApp Element Plug-in for vCenter Server on page 15
Installing a management node on page 18
Modifying vCenter properties for an in-house (dark site) HTTP server on page 25

Related references

Plug-in registration successful but icons do not appear in web client on page 179

Modifying vCenter properties for an in-house (dark site) HTTP server

You must modify the vSphere Web Client properties file if you intend to customize a URL for an in-house (dark site) HTTP server during vCenter Plug-in registration.

Before you begin

You have permissions to download software from the NetApp Support Site.

Steps

1. Follow the appropriate procedure for your installation to modify the `webclient.properties` file to allow vCenter to download from an HTTP server:

Option	Description

vCSA	**a.** SSH into the vCenter Server:
	`Connected to service`
	`* List APIs: "help api list"`
	`* List Plugins: "help pi list"`
	`* Launch BASH: "shell"`
	`Command>`
	b. Enter
	`shell`
	in the command prompt to access root:
	`Command> shell`
	`Shell access is granted to root`
	c. Stop the VMware vSphere Web Client service:
	`service-control --stop vsphere-client`
	`service-control --stop vsphere-ui`
	d. Change the directory:
	`cd /etc/vmware/vsphere-client`
	e. Edit the `webclient.properties` file and add `allowHttp=true`.
	f. Change the directory:
	`cd /etc/vmware/vsphere-ui`
	g. Edit the `webclient.properties` file and add `allowHttp=true`.
	h. Start the VMware vSphere Web Client service:
	`service-control --start vsphere-client`
	`service-control --start vsphere-ui`

Option	Description

Windows	**a.** Change the directory from a command prompt:
	`cd c:\Program Files\VMware\vCenter Server\bin`
	b. Stop the VMware vSphere Web Client service:
	`service-control --stop vsphere-client`
	`service-control --stop vsphere-ui`
	c. Change the directory:
	`cd c:\ProgramData\VMware\vCenterServer\cfg\vsphere-client`
	d. Edit the `webclient.properties` file and add `allowHttp=true`.
	e. Change the directory:
	`cd  c:\ProgramData\VMware\vCenterServer\cfg\vsphere-ui`
	f. Edit the `webclient.properties` file and add `allowHttp=true`.
	g. Change the directory from a command prompt:
	`cd c:\Program Files\VMware\vCenter Server\bin`
	h. Start the VMware vSphere Web Client service:
	`service-control --start vsphere-client`
	`service-control --start vsphere-ui`

Note: After you have completed the registration procedure, you can remove `allowHttp=true` from the files you modified.

2. Reboot vCenter.

Related tasks

Upgrading the vCenter Plug-in on page 15
Installing the NetApp Element Plug-in for vCenter Server on page 15

Upgrading the management node

You can upgrade your management node to the version 11.5 management node after you have successfully upgraded your 10.x version to version 11.0 and later. The vCenter Plug-in 4.3 requires management node 11.3 or later.

Step

1. Choose one of the following management node upgrade options:
   - If you are upgrading from management node 11.3, download and deploy the management node ISO for Element software 11.5 or NetApp HCI 1.7:
     Upgrading the management node to version 11.5 from version 11.3
   - If you are upgrading from management node 11.0 or 11.1, download and deploy the management node ISO for Element software 11.5 or NetApp HCI 1.7:
     Upgrading the management node to version 11.5 from version 11.0 or 11.1
   - If you are upgrading from a management node version 10.x, you must migrate your management node settings to a new management node 11.1 VM before you can update to management node 11.3 and later:
     Migrating from management node 10.x to version 11.x

Migrating from management node 10.x to version 11.x

If you have a management node at version 10.x, you cannot upgrade from 10.x to 11.x. You can instead use this migration procedure to copy over the configuration from 10.x to a newly deployed 11.1 management node. If your management node is currently at 11.0 or higher, you should skip this procedure. You need management node 11.0 or 11.1 and the latest HealthTools to upgrade Element software from 10.3 + through 11.x.

Steps

1. From the VMware vSphere interface, deploy the management node 11.1 OVA and power it on.
2. Open the management node VM console, which brings up the terminal user interface (TUI). Use the TUI to create a new administrator ID and assign a password.
3. In the management node TUI, log in to the management node with the new ID and password and validate that it works.
4. From the vCenter or management node TUI, get the management node 11.1 IP address and browse to the IP address on port 9443 to open the management node UI.

   ```
   https://<mNode 11.1 IP address>:9443
   ```

5. In vSphere, select **NetApp Element Configuration** > **mNode Settings**. In older versions, the top-level menu is **NetApp SolidFire Configuration**.
6. Click **Actions** > **Clear**.
7. To confirm, click **Yes**. The **mNode Status** field should report **Not Configured**.

   Note: When you go to the **mNode Settings** tab for the first time, the **mNode Status** field might display as **Not Configured** instead of the expected **UP**; you might not be able to choose **Actions** > **Clear**. Refresh the browser. The **mNode Status** field will eventually display **UP**.

8. Log out of vSphere.

9. In a web browser, open the management node registration utility (`https://<mNode 11.1 IP address>:9443`) and select **QoSSIOC Service Management**.
10. Set the new QoSSIOC password.

    Note: The default password is `solidfire`. This password is required to set the new password.
11. Click the **vCenter Plug-in Registration** tab.
12. Select **Update Plug-in**.
13. Enter required values. When you are finished, click **UPDATE**.
14. Log in to vSphere and select **NetApp Element Configuration** > **mNode Settings**.
15. Click **Actions** > **Configure**.
16. Provide the management node IP address, management node user ID (the user name is `admin`), password that you set on the **QoSSIOC Service Management** tab of the registration utility, and vCenter user ID and password.

    In vSphere, the **mNode Settings** tab should display the **mNode Status** as **UP**, which indicates management node 11.1 is registered to vCenter.
17. From the management node registration utility (`https://<mNode 11.1 IP address>:9443`), restart the SIOC service from **QoSSIOC Service Management**.
18. Wait for one minute and check the **NetApp Element Configuration** > **mNode Settings** tab. This should display the **mNode Status** as **UP**.

    If status is **DOWN**, check the permissions for `/sf/packages/sioc/app.properties`. The file should have read, write, and execute permissions for the file owner. The correct permissions should appear as follows: `-rwx------`
19. After the SIOC process starts and vCenter displays **mNode Status** as **UP**, check the logs for the `sf-hci-nma` service on the management node. There should be no error messages.
20. (For management node 11.1 only) SSH into the management node version 11.1 with root privileges and start the NMA service with the following commands:

    ```
    # systemctl enable /sf/packages/nma/systemd/sf-hci-nma.service
    ```

    ```
    # systemctl start sf-hci-nma
    ```
21. Perform actions from vCenter to remove a drive, add a drive or reboot nodes. This triggers storage alerts which should be reported in vCenter. If this is working, NMA system alerts are functioning as expected.
22. If ONTAP Select is configured in vCenter, configure ONTAP Select alerts in NMA by copying the `.ots.properties` file from the previous management node to the management node version 11.1 `/sf/packages/nma/conf/.ots.properties` file, and restart the NMA service using the following command: `systemctl restart sf-hci-nma`.
23. Verify that ONTAP Select is working by viewing the logs with the following command: `journalctl -f | grep -i ots`.
24. Configure AIQ by doing the following:

    a. SSH in to the management node version 11.1 and go to the `/sf/packages/collector` directory.

b. Run the following command:

```
sudo ./manage-collector.py --set-username netapp --set-password --
set-mvip <MVIP>
```

c. Enter the management node UI password when prompted.

d. Run the following commands:

```
./manage-collector.py --get-all
```

```
sudo systemctl restart sfcollector
```

e. Verify `sfcollector` logs to confirm it is working.

25. In vSphere, the **NetApp Element Configuration** > **mNode Settings** tab should display the **mNode Status** as **UP**

26. Verify NMA is reporting system alerts and ONTAP Select alerts.

27. If everything is working as expected, shut down and delete management node 10.x VM.

Upgrading the management node to version 11.5 from version 11.3

You can perform an in-place upgrade of the management node from 11.3 to version 11.5 without needing to provision a new management node virtual machine.

Before you begin

- Storage nodes are running Element 11.3 or later.

  Note: Use the latest HealthTools to upgrade Element software.

- The management node you are intending to upgrade is version 11.3 and uses IPv4 networking. The management node 11.5 does not support IPv6.

  Tip: To check the version of your management node, log in to your management node and view the Element version number in the login banner.

- You have updated your management services bundle to the latest version (2.1.368 or later) using one of these options:

  - Hybrid Cloud Control (HCC): `https://<ManagementNodeIP>`

    Note: HCC is available with management services bundle 2.1.326. For NetApp HCI, the HCC option is only available if you have performed the one-time upgrade from NetApp Deployment Engine: `https://<StorageNodeMIP>:442/nde`

  - Management node API: `https://<ManagementNodeIP>/mnode`

    Note: After authenticating, use the mNode API `PUT /services/update/latest` to update services.

- You have configured an additional network adapter (if required) using the instructions for configuring a storage NIC (eth1) in the management node user guide your product.

  Note: Persistent volumes might require an additional network adapter if eth0 is not able to be routed to the SVIP. Configure a new network adapter on the iSCSI storage network to allow the configuration of persistent volumes.

- You have logged in to the management node virtual machine using SSH or console access.

- You have downloaded the management node ISO for *NetApp HCI* or *Element software* from the NetApp Support Site to the management node virtual machine.

  Note: The name of the ISO is similar to `solidfire-fdva-sodium-patch5-11.5.0.xxxx.iso`

- You have checked the integrity of the download by running md5sum on the downloaded file and compared the output to what is available on NetApp Support Site for *NetApp HCI* or *Element software*, as in the following example:

```
sudo md5sum -b <path to iso>/solidfire-fdva-sodium-
patch5-11.5.0.xxxx.iso
```

Steps

1. Mount the management node ISO image and copy the contents to the file system using the following commands:

```
sudo mkdir -p /upgrade
```

```
sudo mount solidfire-fdva-sodium-patch5-11.5.0.xxxx.iso /mnt
```

```
sudo cp -r /mnt/* /upgrade
```

2. Change to the home directory, and unmount the ISO file from `/mnt`:

```
sudo umount /mnt
```

3. Delete the ISO to conserve space on the management node:

```
sudo rm <path to iso>/solidfire-fdva-sodium-patch5-11.5.0.xxxx.iso
```

4. On an 11.3 (11.3.0.14235) management node, run the following command to upgrade your management node OS version. The script retains all necessary configuration files after the upgrade, such as Active IQ collector and proxy settings.

```
sudo /sf/rtfi/bin/sfrtfi_inplace file:///upgrade/casper/
filesystem.squashfs sf_upgrade=1
```

   The management node reboots with a new OS after the upgrade process completes.

5. On the 11.5 management node, run the `redeploy-mnode` script to retain previous management services configuration settings:

   Note: The script retains previous management services configuration, including configuration from the Active IQ collector service, controllers (vCenters), or proxy, depending on your settings.

```
sudo /sf/packages/mnode/redeploy-mnode -mu <mnode user>
```

Related tasks

Upgrading the vCenter Plug-in on page 15
Updating management services with NetApp Hybrid Cloud Control on page 42
Configuring a storage NIC (eth1) on page 35

Upgrading the management node to version 11.5 from version 11.0 or 11.1

You can perform an in-place upgrade of the management node from 11.0 or 11.1 to version 11.5 without needing to provision a new management node virtual machine.

Before you begin

- Storage nodes are running Element 11.3 or later.

  Note: Use the latest HealthTools to upgrade Element software.

- The management node you are intending to upgrade is version 11.0 or 11.1 and uses IPv4 networking. The management node 11.5 does not support IPv6.

  Tip: To check the version of your management node, log in to your management node and view the Element version number in the login banner.

  Note: For management node 11.0, the VM memory needs to be manually increased to 12GB.

- You have configured an additional network adapter (if required) using the instructions for configuring a storage NIC (eth1) in the management node user guide your product.

  Note: Persistent volumes might require an additional network adapter if eth0 is not able to be routed to the SVIP. Configure a new network adapter on the iSCSI storage network to allow the configuration of persistent volumes.

- You have logged in to the management node virtual machine using SSH or console access.
- You have downloaded the management node ISO for *NetApp HCI* or *Element software* from the NetApp Support Site to the management node virtual machine.

  Note: The name of the ISO is similar to `solidfire-fdva-sodium-patch5-11.5.0.xxxx.iso`

- You have checked the integrity of the download by running md5sum on the downloaded file and compared the output to what is available on NetApp Support Site for *NetApp HCI* or *Element software*, as in the following example:

```
sudo md5sum -b <path to iso>/solidfire-fdva-sodium-
patch5-11.5.0.xxxx.iso
```

Steps

1. Mount the management node ISO image and copy the contents to the file system using the following commands:

```
sudo mkdir -p /upgrade
```

```
sudo mount solidfire-fdva-sodium-patch5-11.5.0.xxxx.iso /mnt
```

```
sudo cp -r /mnt/* /upgrade
```

2. Change to the home directory, and unmount the ISO file from `/mnt`:

```
sudo umount /mnt
```

3. Delete the ISO to conserve space on the management node:

```
sudo rm <path to iso>/solidfire-fdva-sodium-patch5-11.5.0.xxxx.iso
```

4. Run one of the following scripts with options to upgrade your management node OS version. Only run the script that is appropriate for your version. Each script retains all necessary configuration files after the upgrade, such as Active IQ collector and proxy settings.

   - On an 11.1 (11.1.0.73) management node, run the following command:

```
sudo /sf/rtfi/bin/sfrtfi_inplace file:///upgrade/casper/
filesystem.squashfs sf_upgrade=1 sf_keep_paths="/sf/packages/
solidfire-sioc-4.2.3.2288 /sf/packages/solidfire-nma-1.4.10/
conf /sf/packages/sioc /sf/packages/nma"
```

   - On an 11.1 (11.1.0.72) management node, run the following command:

```
sudo /sf/rtfi/bin/sfrtfi_inplace file:///upgrade/casper/
filesystem.squashfs sf_upgrade=1 sf_keep_paths="/sf/packages/
solidfire-sioc-4.2.1.2281 /sf/packages/solidfire-nma-1.4.10/
conf /sf/packages/sioc /sf/packages/nma"
```

   - On an 11.0 (11.0.0.781) management node, run the following command:

```
sudo /sf/rtfi/bin/sfrtfi_inplace file:///upgrade/casper/
filesystem.squashfs sf_upgrade=1 sf_keep_paths="/sf/packages/
solidfire-sioc-4.2.0.2253 /sf/packages/solidfire-nma-1.4.8/
conf /sf/packages/sioc /sf/packages/nma"
```

   The management node reboots with a new OS after the upgrade process completes.

5. On the 11.5 management node, run the `upgrade-mnode` script to retain previous configuration settings.

   Note: If you are migrating from an 11.0 or 11.1 management node, the script copies the Active IQ collector to the new configuration format.

   - For a single storage cluster managed by an existing management node 11.0 or 11.1 with persistent volumes:

```
sudo /sf/packages/mnode/upgrade-mnode -mu <mnode user> -pv <true -
persistent volume> -pva <persistent volume account name - storage
volume account>
```

   - For a single storage cluster managed by an existing management node 11.0 or 11.1 with no persistent volumes:

```
sudo /sf/packages/mnode/upgrade-mnode -mu <mnode user>
```

   - For multiple storage clusters managed by an existing management node 11.0 or 11.1 with persistent volumes:

```
sudo /sf/packages/mnode/upgrade-mnode -mu <mnode user> -pv <true -
persistent volume> -pva <persistent volume account name - storage
volume account> -pvm <persistent volumes mvip>
```

- For multiple storage clusters managed by an existing management node 11.0 or 11.1 with no persistent volumes (`-pvm` flag is just to provide one of the cluster's MVIP addresses):

```
sudo /sf/packages/mnode/upgrade-mnode -mu <mnode user> -pvm <mvip
for persistent volumes>
```

6. (For all NetApp HCI installations and SolidFire stand-alone storage installations with NetApp Element Plug-in for vCenter Server) Update the vCenter Plug-in on the 11.5 management node:

   a. Log out of the vSphere Web Client.

      Note: The web client will not recognize updates made during this process to your vCenter Plug-in if you do not log out.

   b. Browse to the registration utility (`https://<ManagementNodeIP>:9443`).

   c. Click the **vCenter Plug-in Registration** tab.

   d. Within **Manage vCenter Plug-in**, select **Update Plug-in**.

   e. Update the vCenter address, vCenter administrator user name, and vCenter administrator password.

   f. Click **Update**.

   g. Log in to the vSphere Web Client and verify that the plug-in information has been updated by browsing to **Home > NetApp Element Configuration > About**.

      Note: Logging into vSphere Web Client after updating registration installs the new plug-in updates and creates a new database.

      You should see the following version details:

      - NetApp Element Plug-in Version: 4.3.0
      - NetApp Element Plug-in Build Number: 233

7. Use the mNode API to add assets:

   a. Using a browser, go to the storage MVIP and log in.

      This action causes certificate to be accepted for the next step.

   b. Using a browser, go to `https://<ManagementNodeIP>/mnode`.

   c. Add a vCenter controller asset to the management node known assets for HCI monitoring (NetApp HCI installations only) and Hybrid Cloud Control (HCC):

   d. Click **Authorize** and enter your MVIP user name and password credentials. Close the pop-up window.

   e. Run **GET /assets** to pull the base asset ID needed to add the vCenter/controller asset.

   f. Run **POST /assets/{ASSET_ID}/controllers** to add a controller asset with vCenter credentials.

   g. For NetApp HCI or to access cloud services options in HCC, add a compute asset to the management node known assets:

      Attention: You must perform this step or Cloud Services options will not be available from HCC.

   h. Click **Authorize** and enter your MVIP user name and password credentials. Close the pop-up window.

i. Run **GET /assets** to pull the base asset ID needed to add the compute asset.

j. Run **POST/assets/{asset_id}/compute-nodes** to add a compute asset with credentials for the compute asset. The type is `ESXi Host`.

Related tasks

Upgrading the vCenter Plug-in on page 15
Updating management services with NetApp Hybrid Cloud Control on page 42
Configuring a storage NIC (eth1) on page 35

Configuring a storage NIC (eth1)

If you are using a second NIC for storage, you can SSH in to the management node or use the vCenter console and run a curl command to set up the network for eth1.

Before you begin

- You know your eth0 configuration details.
- Your cluster version is running NetApp Element software 11.3 or later.
- You have deployed a management node 11.3 or later.

Steps

1. Open an SSH or vCenter console.
2. Replace the values in the following command template (represented by `$`) for each of the required parameters for eth0 and eth1:

   Note: The `cluster` object in the following template is optional and can be used for management node host name renaming. The `- -insecure` and the `-k` options should not be used in production environments.

```
curl -u $mnode-username:$mnode-password --insecure -X POST \
  https://$mnode_management_IP:442/json-rpc/10.0 \
  -H 'Content-Type: application/json' \
  -H 'cache-control: no-cache' \
  -d '  {
        "params": {
                "network": {
                        "eth0": {
                                "address": "$eth0_ip_mnode_management_IP",
                                "dns-nameservers": "$dns_ip_or_hostname",
                                "netmask": "$eth0_net_mask",
                                "gateway": "$gateway_IP",
                                "gatewayV6": ""
                        },
                        "eth1": {
                                "address": "$eth1_ip",
                                "netmask": "$eth1_netmask",
                                "status": "Up",
                                "method": "static",
                                "mtu": "9000"
                        }
      },
                "cluster": {
                        "name": "$desired_mNode_vm_hostname"
                }
        },
        "method": "SetConfig"
}
'
```

3. Run the command.

Persistent volumes

Persistent volumes allow management node configuration data to be stored on a specified storage cluster, rather than locally with a VM, so that data can be preserved in the event of management node loss or removal. Persistent volumes are an optional but recommended management node configuration.

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

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

Attention: Persistent volumes are assigned to a new account that is also created during installation or upgrade. After you created persistent volumes, you must not modify or delete the volumes and their associated account.

Updating the vCenter Plug-in registration for vCSA

You can update the vCenter Plug-in for a vCenter Server Virtual Appliance (vCSA) using the registration user interface. The registration UI is available from the IP address for your cluster's management node.

Before you begin

- You have logged out of the vSphere Web Client.

  Note: The web client will not recognize updates made during this process to your vCenter Plug-in if you do not log out.

- You have previously registered the vCenter Plug-in 3.0.1 or later on vCenter Server.
- You have vCenter Administrator role privileges to register a plug-in.
- You have deployed a management node OVA running Element software 11.3 or later.
- You have upgraded your management node to 11.3 or later.
- Your management node is powered on with its IP address or DHCP address configured.
- You are using an SSH client or web browser (Chrome 56 or later or Firefox 52 or later).
- Your firewall rules allow open network communication between the vCenter and the storage cluster MVIP on TCP ports 443, 8443, and 9443. Port 9443 is used for registration and can be closed after registration is complete. If you have enabled virtual volumes functionality on the cluster, ensure TCP port 8444 is also open for VASA provider access.

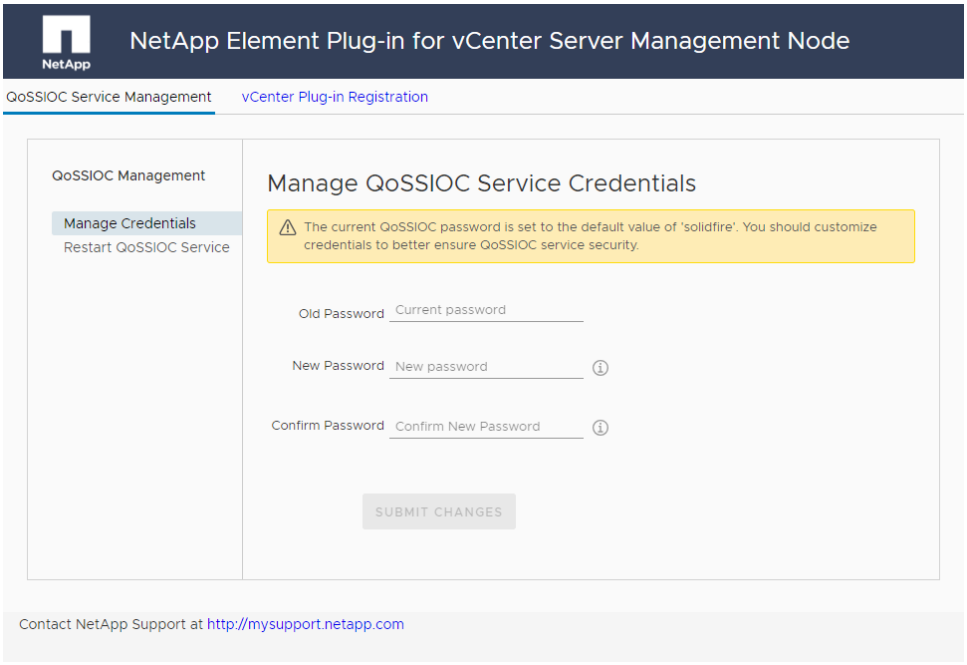
About this task

You must change your registration for the vCenter Plug-in on every vCenter Server where you need to use the plug-in.

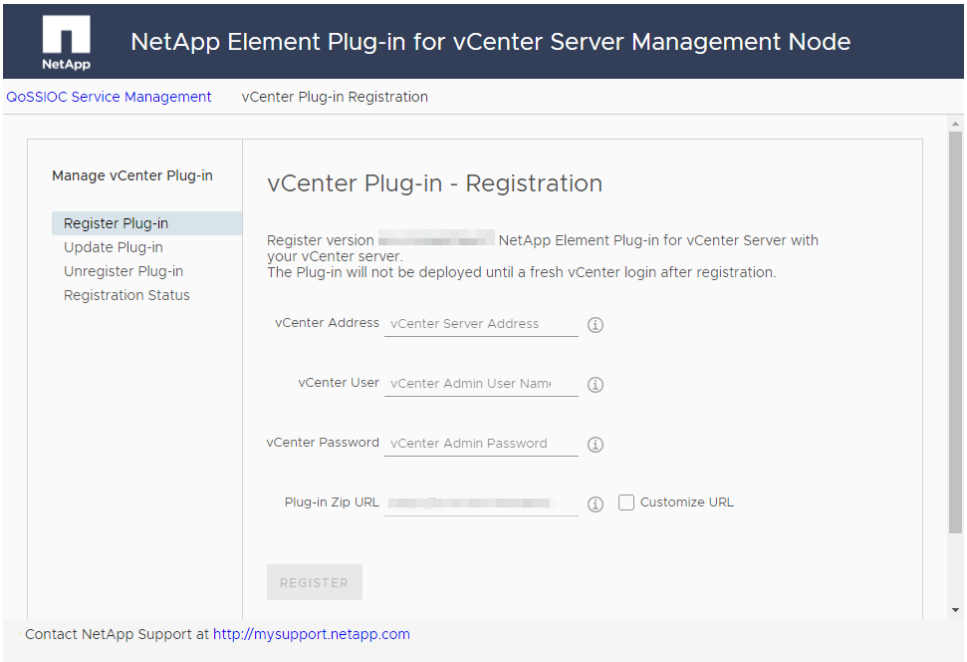
Steps

1. Enter the IP address for your management node in a browser, including the TCP port for registration: `https://[management node IP]:9443`.

   The registration UI displays the **Manage QoSSIOC Service Credentials** page for the plug-in.



2. Click **vCenter Plug-in Registration**.



3. Click **Registration Status**, complete the necessary fields, and click **Check Status** to verify that the vCenter Plug-in is already registered and the version number of the current installation.

4. Click **Update Plug-in**.
5. Confirm or update the following information:
   - The IPv4 address or the FQDN of the vCenter service on which you will register your plug-in.
   - The vCenter Administrator user name.

     Note: The user name and password credentials you enter must be for a user with vCenter Administrator role privileges.

   - The vCenter Administrator password.
   - (For in-house servers/dark sites) A custom URL for the plug-in ZIP.

     Note: Most installations use the default path. You can click **Custom URL** to customize the URL if you are using an HTTP or HTTPS server (dark site) or have modified the ZIP file name or network settings. For additional configuration steps if you intend to customize a URL, see plug-in documentation about modifying vCenter properties for an in-house (dark site) HTTP server.
6. Click **Update**.
7. Optional: Click **Registration Status**, complete the necessary fields, and click **Check Status** to verify that the updated version of the plug-in is registered on the vCenter.
8. (vCenter Plug-in 3.0 and earlier only) If you are updating from any vCenter Plug-in 3.0 or earlier, type the following commands from an SSH session on the vCSA as root account:

   ```
   # rm -rf /storage/vsphere-client/netapp-solidfire
   # reboot
   ```

   Note: This action removes the version 3.0.0 database prior to the installation of the new vCenter Plug-in in the vSphere Web Client. After completing the registration process, you need to add clusters again using the NetApp Element Configuration extension point.
9. Log in to the vSphere Web Client as a vCenter Administrator.

   Note: This action installs the new plug-in updates and creates a new database. If the vCenter Plug-in icons are not visible from the vSphere main page, see documentation about troubleshooting the plug-in.
10. Verify the version change in the **About** tab in the **NetApp Element Configuration** extension point.

    Note: The vCenter Plug-in contains online Help content. To ensure that your Help contains the latest content, clear your browser cache after upgrading your plug-in.

Related tasks

Upgrading the vCenter Plug-in on page 15
Installing a management node on page 18
Modifying vCenter properties for an in-house (dark site) HTTP server on page 25

Related references

Plug-in registration successful but icons do not appear in web client on page 179

Updating the vCenter plug-in registration for Windows vCenter

You can update the vCenter Plug-in for a vCenter Server on Windows from the registration user interface. The registration UI is available from the IP address for your cluster's management node.

Before you begin

- You have logged out of the vSphere Web Client.

  Note: The web client will not recognize updates made during this process to your vCenter Plug-in if you do not log out.

- You have previously registered the vCenter Plug-in 3.0.1 or later on vCenter Server.
- You have vCenter Administrator role privileges to register a plug-in.
- You have deployed a management node OVA running Element software 11.3 or later.
- You have upgraded your management node to 11.3 or later.
- Your management node is powered on with its IP address or DHCP address configured.
- You are using an SSH client or web browser (Chrome 56 or later or Firefox 52 or later).
- Your firewall rules allow open network communication between the vCenter and the storage cluster MVIP on TCP ports 443, 8443, and 9443. Port 9443 is used for registration and can be closed after registration is complete. If you have enabled virtual volumes functionality on the cluster, ensure TCP port 8444 is also open for VASA provider access.

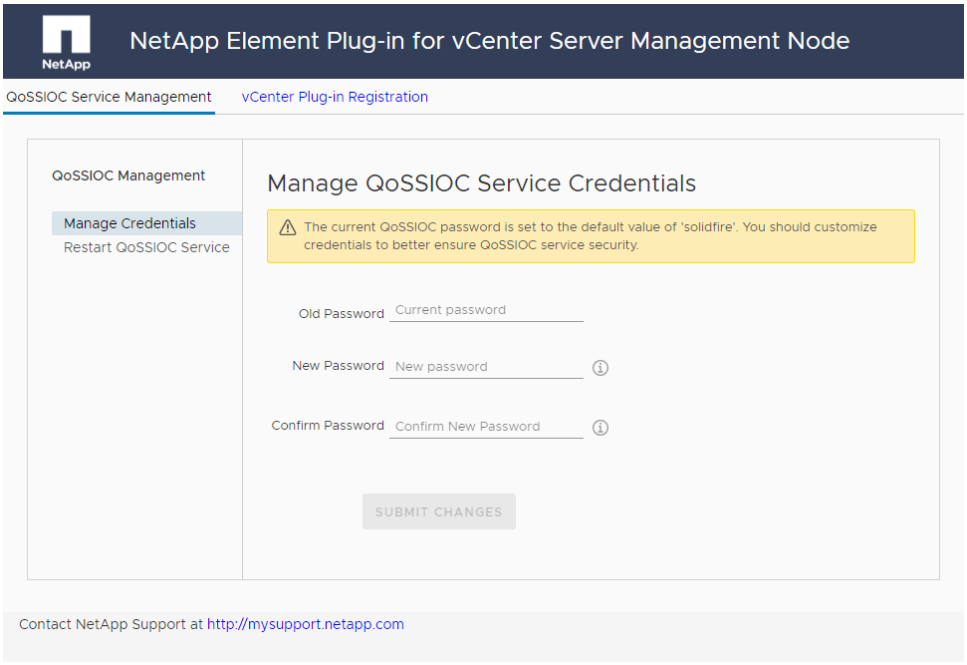
About this task

You must change your registration for the vCenter Plug-in on every vCenter Server where you need to use the plug-in.

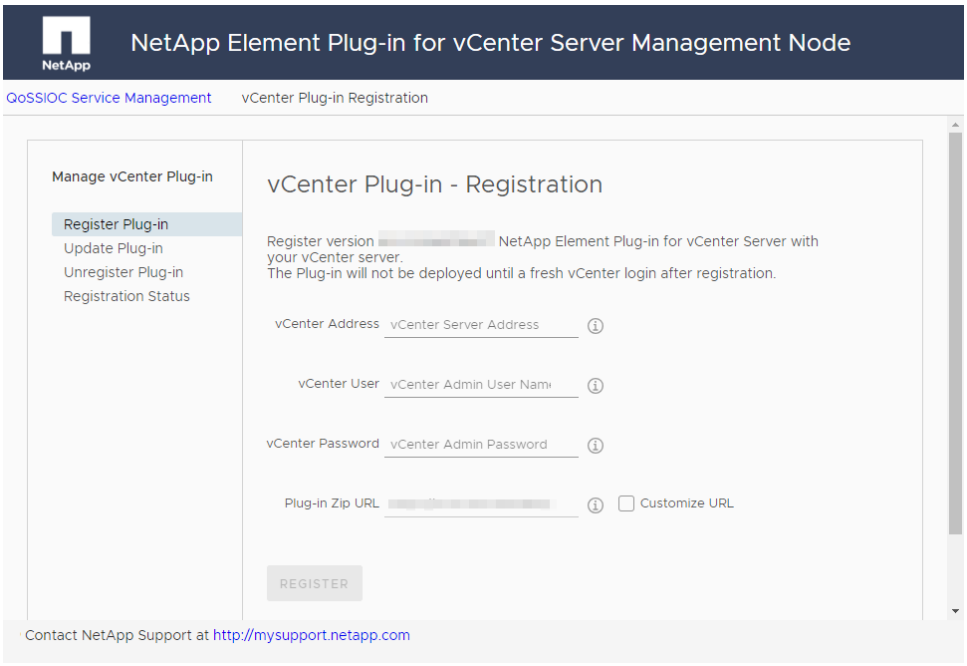
Steps

1. Enter the IP address for your management node in a browser, including the TCP port for registration: `https://[management node IP]:9443`.

   The registration UI displays the **Manage QoSSIOC Service Credentials** page for the plug-in.



2. Click **vCenter Plug-in Registration**.



3. Click **Registration Status**, complete the necessary fields, and click **Check Status** to verify that the vCenter Plug-in is already registered and the version number of the current installation.
4. Click **Update Plug-in**.
5. Confirm or update the following information:
   - The IPv4 address or the FQDN of the vCenter service on which you will register your plug-in.
   - The vCenter Administrator user name.

Note: The user name and password credentials you enter must be for a user with vCenter Administrator role privileges.

- The vCenter Administrator password.
- (For in-house servers/dark sites) A custom URL for the plug-in ZIP.

  Note: Most installations use the default path. You can click **Custom URL** to customize the URL if you are using an HTTP or HTTPS server (dark site) or have modified the ZIP file name or network settings. For additional configuration steps if you intend to customize a URL, see plug-in documentation about modifying vCenter properties for an in-house (dark site) HTTP server.

6. Click **Update**.
7. For vCenter Plug-in 3.0 and earlier only, complete the following steps:
   a. Delete the folder **serenity** and all its content. For Windows vCenter Server 6.x, the path is the following: `C:\ProgramData\VMware\vCenterServer\data\vSphere Web Client \SerenityDB\serenity`
   b. Purge the deleted content from the Recycle Bin.
   c. Reboot vCenter.
8. Log in to the vSphere Web Client as a vCenter Administrator.

   Note: This action completes the installation in the vSphere Web Client. If the vCenter Plug-in icons are not visible from the vSphere main page, see documentation about troubleshooting the plug-in.

9. Verify the version change in the **About** tab in the **NetApp Element Configuration** extension point.

   Note: The vCenter Plug-in contains online Help content. To ensure that your Help contains the latest content, clear your browser cache after upgrading your plug-in.

Related tasks

Upgrading the vCenter Plug-in on page 15
Installing a management node on page 18
Modifying vCenter properties for an in-house (dark site) HTTP server on page 25

Related references

Plug-in registration successful but icons do not appear in web client on page 179

Updating management services

You can update your management services to the latest version after you have installed management node 11.3 or later. Beginning with the management node 11.3 release, the QoSSIOC service for the plug-in can be updated independently from Element software releases.

Step

1. Choose one of the following management services update options:
   - Update management services from Hybrid Cloud Control (HCC):
     Updating management services with NetApp Hybrid Cloud Control

- Update management services using the management node API:
  Updating management services using mNode API

Updating management services with NetApp Hybrid Cloud Control

Using NetApp Hybrid Cloud Control, you can update your NetApp management services. Management service bundles provide functionality and fixes to your installation outside of major releases.

Before you begin

- You are running management node 11.3 or later.
- You have upgraded your management services to at least version 2.1.326. NetApp Hybrid Cloud Control upgrades is not available in earlier service bundle versions.

About this task

For a list of available services for each service bundle version, see the *Management Services Release Notes*.

Steps

1. Open a web browser and browse to the IP address of the management node:

   ```
   https://<ManagementNodeIP>
   ```

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Click **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select the **Management Services** tab.

   The **Management Services** tab shows the current and available versions of management services software.

   Note: If your installation cannot access the internet, only the current software version is shown.

5. Do one of the following:

Option	Description
Your installation can access the internet	If a management services upgrade is available, click **Begin Upgrade**.
Your installation cannot access the internet	**a.** Follow the instructions on the page to download and save a management services upgrade package on your computer. **b.** Click **Browse** to locate the package you saved and upload it.

The upgrade begins, and you can see the upgrade status on this page.

Related information

Management services release notes

Updating management services using mNode API

Users should perform management services updates from the NetApp Hybrid Cloud Control page. You alternately can manually update management services using the REST API UI from the

management node. Management services updates are available as service bundles from an online software repository.

Before you begin

- You have internet access.
- You have deployed a NetApp Element software management node 11.3 or later.
- Your cluster version is running NetApp Element software 11.3 or later.

About this task

This procedure describes the manual update of management services using the management node API. Management services include the SIOC service for the Element Plug-in for vCenter, the Active IQ collector service, the NetApp HCI monitoring service (for NetApp HCI installations only) and additional services. Updates to non-service-based components of the management node are provided as updated images (OVA or ISO) and cannot be updated using this procedure.

Steps

1. Open the REST API UI on the management node: `https://[management node IP]/mnode`
2. Click **Authorize** and complete the following:
   a. Enter the cluster user name and password.
   b. Enter the client ID as `mnode-client` if the value is not already populated.
   c. Copy the token URL string and paste it into another browser tab to initiate a token request.
   d. Click **Authorize** to begin a session.
3. (Optional) Confirm available versions of management node services: `GET /services/versions`
4. (Optional) Get detailed information about the latest version: `GET /services/versions/latest`
5. (Optional) Get detailed information about a specific version: `GET /services/versions/{version}/info`
6. Perform one of the following management services update options:

Option	Description
`PUT /services/update/latest`	Run this command to update to the most recent version of management node services.
`PUT /services/update/{version}`	Run this command to update to a specific version of management node services.

7. Use `GET/services/update/status` to monitor the status of the update.

   A successful update returns a result similar to the following example:

```
{
    "current_version": "2.1.346",
    "details": "Updated to version 2.1.346",
    "status": "success"
}
```

Related tasks

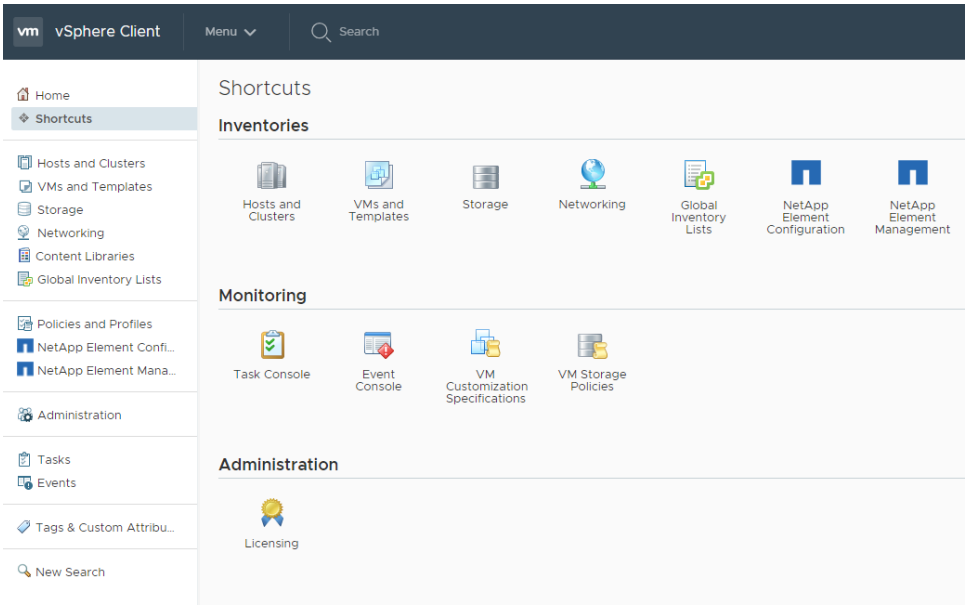
Updating management services with NetApp Hybrid Cloud Control on page 42

Related information

Management services release notes

Accessing the plug-in after installation

After successful installation, NetApp Element Configuration and Management extension points appear in the **Shortcuts** tab of the vSphere Web Client and in the side panel.



Note: If the vCenter Plug-in icons are not visible, see documentation about troubleshooting the plug-in.

Related tasks

Upgrading the vCenter Plug-in on page 15
Installing the NetApp Element Plug-in for vCenter Server on page 15

Related references

Plug-in registration successful but icons do not appear in web client on page 179

How to use the NetApp Element Plug-in for vCenter Server

The NetApp Element Plug-in for vCenter Server enables you to configure, manage, and monitor NetApp Element clusters in the VMware vSphere Web Client. You can make cluster-wide changes using the configuration and management extension points.

The NetApp Element Configuration extension point allows you to add and manage clusters, assign vCenter Servers for Linked Mode, and configure management node settings for QoSSIOC. The NetApp Element Management extension point gives you a comparable monitoring and management interface to the Element UI for central, cluster-wide control of your storage system.

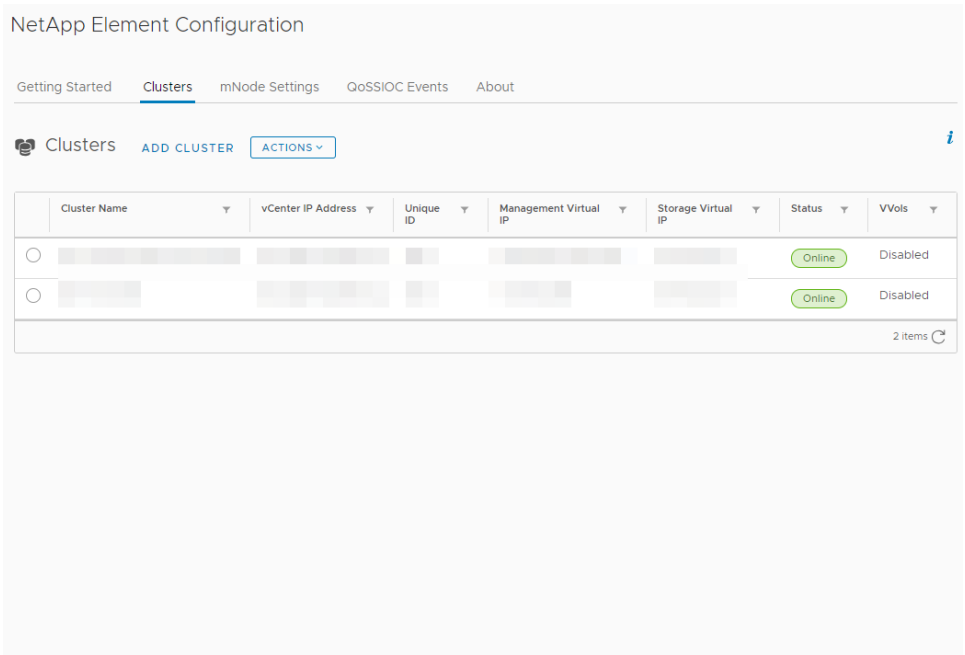
Related references

NetApp Element Configuration Extension Point on page 45
NetApp Element Management Extension Point on page 46

NetApp Element Configuration Extension Point

The NetApp Element Configuration extension point allows you to add and manage clusters, assign vCenter Servers for Linked Mode, and configure management node settings for QoSSIOC.

Note: Your vSphere Web Client might differ slightly from what is shown in the following screen depending on the version of vSphere installed.



The following tabs are available from the NetApp Element Configuration extension point:

Getting Started

Introduces the extension points for the plug-in and the actions that can be performed. You can hide Getting Started pages from each page or restore them from the **About** tab in the NetApp Element Configuration extension point.

Clusters

Manage the NetApp Element clusters controlled by the plug-in. You can also enable, disable, or configure cluster-specific features.

mNode Settings

Configure the management node settings for the QoSSIOC service.

QoSSIOC Events

Displays information about all detected QoSSIOC events.

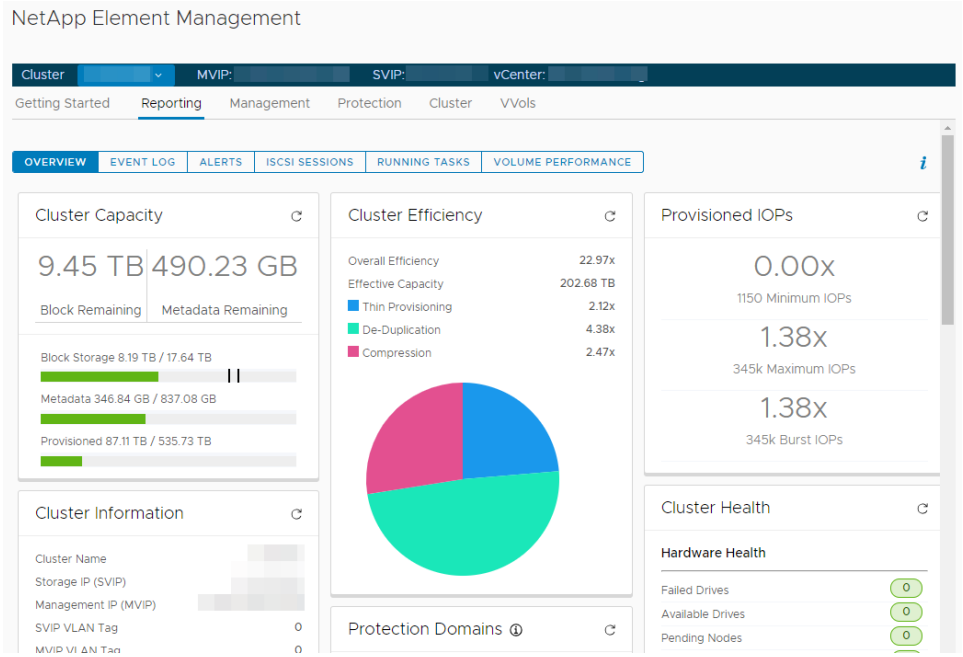
About

Displays plug-in version information and provides a service bundle download option.

NetApp Element Management Extension Point

The NetApp Element Management extension point gives you a comparable monitoring and management interface to the Element UI for central, cluster-wide control of your storage system.

Note: Your vSphere Web Client might differ slightly from what is shown in the following screen depending on the version of vSphere installed.



The cluster navigation bar allows you to quickly switch between clusters that have been added to the plug-in:

Cluster

If two or more clusters are added, ensure that the cluster you intend to use for management tasks is selected in the navigation bar. Select other added clusters from the drop-down list.

MVIP

The management virtual IP address of the selected cluster.

SVIP

The storage virtual IP address of the selected cluster.

vCenter

The vCenter Server which the selected cluster can access. The cluster is assigned access to a vCenter Server when the cluster is added to the plug-in.

The following tabs are available from the NetApp Element Management extension point:

Getting Started

Introduces the extension points for the plug-in and the actions that can be performed. You can hide Getting Started pages from each page or restore them from the **About** tab in the NetApp Element Configuration extension point.

Reporting

Displays information about cluster components and provides a cluster performance overview. You can also find information about events, alerts, iSCSI sessions, running tasks, and volume performance from the tab.

Management

Create and manage datastores, volumes, user accounts, access groups, and initiators. You can also perform backup operations, clones, and snapshots. QoS policies are available to be created and managed using NetApp Element software 10 or later.

Protection

Manage individual and group snapshots. You can also create schedules for snapshot creation, pair clusters for real-time replication, and manage volume pairs.

Cluster

Add and manage drives and nodes. You can also create and manage VLANs.

VVols

Manage virtual volumes and their associated storage containers, protocol endpoints, and bindings.

vCenter Linked Mode

You can use the NetApp Element Plug-in for vCenter Server to manage cluster resources from other vCenter Servers using vCenter Linked Mode. You can log into any vCenter Server that is part of a Linked Mode group and manage the resources owned by other linked vCenter Servers from a single interface.

The plug-in must be registered with each vCenter Server in the Linked Mode environment that will be using the plug-in. You must complete registration with each vCenter Server using the registration utility for the plug-in. You must also log in one time to the vSphere Web Client for each linked vCenter Server. Logging in initiates installation of the plug-in on the web client.

Best Practices: Manage the cluster from the vCenter Server you associate with the cluster during the **Add Cluster** configuration process.

Hosts that use Element software-based storage are exclusive to a particular vCenter Server and are not shared between members of a Linked Mode group. Because of this, any storage management tasks for a cluster are limited to the available hosts within a vCenter Server.

Related tasks

Adding a cluster on page 50

Using object naming best practices when managing multiple clusters

You should use a consistent and descriptive naming convention for all objects associated with the NetApp Element Plug-in for vCenter Server to ensure you can identify cluster components easily. This is particularly important when using multiple vCenter Server installations and multiple clusters.

Managing multiple clusters can become confusing if a descriptive naming convention is not used. For example, if you have `SF-cluster1` and `SF-cluster2` in your inventory, accounts, volumes, and datastores should follow a similar pattern, such as `SF1-account1` and `SF2-vol1`. Using consistent naming makes it easier to identify the cluster you are working with without having to examine the cluster navigation bar in the **NetApp Element Management** extension point. It also minimizes the chances of incorrectly modifying an object in vSphere that serves as the backing for an Element software-based storage object; for example, a datastore that is associated with an Element software-based volume.

Cluster configuration

The **Clusters** tab in the **NetApp Element Configuration** extension point allows you to add clusters that you can then manage from the **NetApp Element Management** extension point. You can also manage existing cluster profiles, enable Virtual Volumes (VVols) on clusters that support the functionality, or shut down a cluster to take it offline.

Adding a cluster

You can add a cluster running Element software using the NetApp Element Configuration extension point. After a connection has been established to the cluster, the cluster can then be managed using the NetApp Element Management extension point.

Before you begin

- At least one cluster must be available and its IP or FQDN address known.
- Current full Cluster Admin user credentials for the cluster.
- Firewall rules allow open network communication between the vCenter and the cluster MVIP on TCP ports 443 and 8443.

  Note: You must add at least one cluster to use NetApp Element Management extension point functions.

About this task

This procedure describes how to add a cluster profile so that the cluster can be managed by the plug-in. You cannot modify cluster administrator credentials using the plug-in. For instructions on changing credentials for a cluster administrator account, see the *NetApp Element Software User Guide*.

Attention: The vSphere HTML5 web client and Flash web client have separate databases that cannot be combined. Clusters added in one client will not be visible in the other. If you intend to use both clients, add your clusters in both.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Click **Add Cluster**.
3. In the **Add Cluster** dialog box, enter the following information:
   - **IP address/FQDN**: Enter the cluster MVIP address.
   - **User ID**: Enter a cluster administrator user name.
   - **Password**: Enter a cluster administrator password.
   - **vCenter Server**: If you have set up a Linked Mode group, select the vCenter Server you want to access the cluster. If you are not using Linked Mode, the current vCenter Server is the default.

     Note: The hosts for a cluster are exclusive to each vCenter Server. Be sure that the vCenter Server you select has access to the intended hosts. You can remove a cluster, reassign it to another vCenter Server, and add it again if you decide later to use different hosts.

4. Click **OK**.

   When the process completes, the cluster appears in the list of available clusters and can be used in the NetApp Element Management extension point.

Viewing cluster details

You can view a brief summary of details for each available cluster.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the check box for the cluster profile you want to review.
3. Click **Actions**.
4. In the resulting menu, select **Details**.

Cluster details

You can view cluster information for all clusters that have been added to the plug-in on the **Clusters** page of the **NetApp Element Configuration** extension point.

Cluster Name
The name for the cluster.

vCenter IP Address
The IP address or FQDN of the vCenter Server to which the cluster is assigned.

Unique ID
Unique ID for the cluster.

Management Virtual IP
The management virtual IP address (MVIP).

Storage Virtual IP
The storage virtual IP address (SVIP).

Status
The status of the cluster.

VVols
The status of the VVols functionality on the cluster.

Individual cluster details

You can view detailed information for an individual cluster when you select it and view its details on the **Clusters** page of the **NetApp Element Configuration** extension point.

Cluster Name
The name for the cluster.

Unique ID
Unique ID for the cluster.

vCenter IP Address
The IP address or FQDN of the vCenter Server to which the cluster is assigned.

Management Virtual IP
The management virtual IP address (MVIP).

MVIP Node ID

The node that holds the master MVIP address.

Storage Virtual IP

The storage virtual IP address (SVIP).

SVIP Node ID

The node holding the master SVIP address.

Element Version

The version of NetApp Element software that the cluster is running.

VASA 2 Status

The status of the VASA Provider on Element cluster.

VASA Provider URL

The URL of the VASA Provider enabled on the Element cluster, when applicable.

Encryption At Rest Status

The status of Encryption at Rest.

Possible values:

- `Enabling`: Encryption at Rest is being enabled.
- `Enabled`: Encryption at Rest is enabled.
- `Disabling`: Encryption at Rest is being disabled.
- `Disabled`: Encryption at Rest is disabled.

Ensemble Nodes

IPs of the nodes that are part of the database ensemble.

Paired With

The names of additional clusters that are paired with the local cluster.

SSH Status

The status of the secure shell. If enabled, the time remaining is displayed.

Editing cluster profiles

You can change the Cluster Admin user name and password for a cluster profile from the NetApp Element Configuration extension point.

About this task

This procedure describes how to change the cluster admin user name and password used by the plug-in. You cannot change the cluster admin credentials from the plug-in. For instructions on changing credentials for a cluster administrator account, see the *NetApp Element Software User Guide*.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the check box for the cluster profile you want to edit.
3. Click **Actions**.
4. In the resulting menu, click **Edit**.

5. In the **Edit Cluster** dialog box, change any of the following:
   - **User ID**: The cluster administrator user name.
   - **Password**: The cluster administrator password.

     Note: You cannot change the IP address or FQDN of a cluster profile after a cluster is added. You also cannot change the assigned Linked Mode vCenter Server for an added cluster. To change the cluster address or associated vCenter Server, you must remove the cluster and add it again.
6. Click **OK**.

Removing a cluster profile

You can remove the cluster profile of a cluster that you no longer want to manage from the vCenter Plug-in using the NetApp Element Configuration extension point.

About this task

If you have set up a Linked Mode group and want to reassign a cluster to another vCenter Server, you can remove the cluster profile and add it again with a different linked vCenter Server IP.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the check box for the cluster profile you want to remove.
3. Click **Actions**.
4. In the resulting menu, click **Remove**.
5. Confirm the action.

Enabling virtual volumes

You must manually enable vSphere Virtual Volumes (VVols) functionality using the NetApp Element Configuration extension point. The Element 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 Element cluster must be connected to an ESXi 6.0 and later environment that is compatible with VVols.
- If you are using Element 11.3 or later, the cluster must be connected to an ESXi 6.0 update 3 or later environment.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select a cluster from the list that you want to enable.
3. Click **Actions**.
4. In the resulting menu, click **Enable VVols**.

Attention: After VVols functionality is enabled, it cannot be disabled. Enabling vSphere Virtual Volumes functionality permanently changes NetApp Element software configuration. You should only enable VVols functionality if your cluster is connected to a VMware ESXi VVols-compatible environment. You can only disable the VVols feature and restore the default settings by returning the cluster to the factory image.

5. Click **Yes** to confirm the Virtual Volumes configuration change.

   Note: When VVols functionality is enabled, the Element 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.

6. Click **Actions** for the selected cluster.
7. In the resulting menu, select **Details**.
8. Copy the VASA Provider URL from the **VASA Provider URL** field. You will use this URL to register the VASA Provider in vCenter.

   Note: See plug-in documentation for additional configuration tasks that are required for vSphere Virtual Volumes functionality.

Related tasks

Configuring vSphere Virtual Volumes (VVols) functionality on page 161

Enabling encryption at rest

You can manually enable encryption at rest (EAR) functionality using the NetApp Element Configuration extension point.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the cluster on which you want to enable encryption at rest.
3. Click **Actions**.
4. In the resulting menu, click **Enable EAR**.
5. Confirm the action.

Disabling encryption at rest

You can manually disable encryption at rest (EAR) functionality using the NetApp Element Configuration extension point.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the check box for the cluster on which you want to disable encryption at rest.
3. Click **Actions**.
4. In the resulting menu, click **Disable EAR**.
5. Confirm the action.

Enabling SSH

You can manually enable a Secure Shell (SSH) session using the NetApp Element Configuration extension point. Enabling SSH allows NetApp technical support engineers access to storage nodes for troubleshooting for the duration you determine.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the cluster for which you want to enable an SSH session.
3. Click **Actions**.
4. In the resulting menu, click **Enable SSH**.
5. Enter a duration for the SSH session to be enabled in hours up to a maximum of 720. You must set a value to proceed.
6. Click **Yes**.

Disabling SSH

You can manually disable Secure Shell (SSH) access to nodes in the storage cluster using the NetApp Element Configuration extension point.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the cluster on which you want to disable SSH access.
3. Click **Actions**.
4. In the resulting menu, click **Disable SSH**.
5. Click **Yes**.

Changing the SSH time limit

You can change the duration of an active Secure Shell (SSH) session using the NetApp Element Configuration extension point.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the cluster with the SSH session you want to change.
3. Click **Actions**.
4. In the resulting menu, click **Change SSH**.

   The dialog box displays the remaining time for the SSH session.
5. Enter a new duration for the SSH session in hours. You must set a value to proceed.
6. Click **Yes**.

Setting protection domain monitoring

You can manually enable protection domain monitoring using the NetApp Element Configuration extension point. You can select a protection domain threshold based on node or chassis domains.

About this task

A chassis domain emphasizes the resiliency of the cluster to withstand a chassis-level failure. A node domain emphasizes a select group of nodes, potentially across chassis. A chassis domain requires more potential capacity resources than a node domain to be resilient to failure. When a protection domain threshold is exceeded, a cluster no longer has sufficient capacity to heal from failure while also maintaining undisrupted data availability.

Note: The selected cluster must be Element 11.0 or later to use protection domain monitoring; otherwise, protection domain functions are not available.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the cluster on which you want to enable protection domain monitoring.
3. Click **Actions**.
4. In the resulting menu, click **Set Protection Domain Monitoring**.
5. In the **Set Protection Domain Monitoring** dialog, select a failure threshold:
   - **Node**: The threshold beyond which a cluster can no longer provide uninterrupted data during hardware failures at node level. The node threshold is the system default.
   - **Chassis**: The threshold beyond which a cluster can no longer provide uninterrupted data during hardware failures at chassis level.
6. Click **OK**.

Shutting down a cluster

You can manually shut down all active nodes in a storage cluster using the NetApp Element Configuration extension point.

Before you begin

You have stopped I/O and disconnected all iSCSI sessions.

Note: If you want to restart rather than shut down the cluster, you can select all nodes from the Cluster page in the NetApp Element Management extension point and perform a restart.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the check box for the cluster you want to shut down.
3. Click **Actions**.
4. In the resulting menu, click **Shutdown**.
5. Confirm the action.

Related tasks

Shutting down cluster nodes on page 157
Restarting cluster nodes on page 157

Related information

Powering off and powering on a NetApp HCI system

Expanding your NetApp HCI

You can manually expand your NetApp HCI infrastructure by adding nodes using NetApp HCI. A link to a NetApp HCI UI for scaling your system is provided from the NetApp Element extension point. Additional links are provided within the NetApp Element Management extension point from the **Getting Started** and **Cluster** pages.

Steps

1. Select **NetApp Element Configuration > Clusters**.
2. Select the cluster you want to change.
3. Click **Actions**.
4. In the resulting menu, click **Expand your NetApp HCI**.

Management node (mNode) settings

The **mNode Settings** tab in the NetApp Element Configuration extension point allows you to configure settings that are used for the QoSSIOC service. After you have configured a valid management node, these settings become the default. You can edit these settings to update credentials or clear these settings for a new management node.

For Linked Mode, the NetApp Element Plug-in for vCenter Server registers all vCenter Servers using the management node settings you provide on a single vCenter Server.

The **mNode Status** field on the settings page displays the following possible values:

- `Up`: QoSSIOC is enabled.
- `Down`: QoSSIOC is not enabled.
- `Not Configured`: QoSSIOC settings have not been configured.
- `Network Down`: vCenter cannot communicate with the QoSSIOC service on the network. The mNode and SIOC service might still be running.

Configuring management node settings for QoSSIOC

You can configure the Element management node (mNode) settings using the NetApp Element Configuration extension point. These configurations are required to enable and use the QoSSIOC service.

Steps

1. Select **NetApp Element Configuration > mNode Settings**.
2. Click **Actions**.
3. In the resulting menu, select **Configure**.
4. In the **Configure mNode Settings** dialog box, enter the following information:
   - **mNode IP Address/FQDN**: The IP address of the management node for the cluster that contains the QoSSIOC service.
   - **mNode Port**: The port address for the management node that contains the QoSSIOC service. The default port is 8443.
   - **mNode User ID**: The user ID for the QoSSIOC service. The QoSSIOC service default user id is *admin*. For NetApp HCI, the user ID is the same one entered during installation using the NetApp Deployment Engine.
   - **mNode Password**: The password for the Element QoSSIOC service. The QoSSIOC service default password is *solidfire*. If you have not created a custom password, you can create one from the Registration UI (`https://[management node IP]:9443`).
   - **vCenter User ID**: The user name for the vCenter admin with full Administrator role privileges.
   - **vCenter Password**: The password for the vCenter admin with full Administrator role privileges.
5. Click **OK**.

The **mNode Status** field displays UP when the plug-in can successfully communicate with the service.

Note: After you have configured a valid management node, these settings become the default. The mNode settings revert to the last known valid mNode settings until you provide settings for another valid management node. You must clear the settings for the configured management node before setting the credentials for a new management node.

Related information

KB article: vCenter Plugin (VCP) mnode settings credentials are no longer valid after management node redeployment or VCP re-registration on an HCI cluster

Editing management node settings for QoSSIOC

You can change the management node (mNode) and vCenter credentials of an active Element management node using the NetApp Element Configuration extension point.

Steps

1. Select **NetApp Element Configuration** > **mNode Settings**.
2. Click **Actions**.
3. In the resulting menu, select **Edit**.
4. In the **Edit mNode Settings** dialog box, change any of the following:
   - **mNode User ID**: The user ID for the QoSSIOC service. The QoSSIOC service default user id is *admin*. For NetApp HCI, the user ID is the same one entered during installation using the NetApp Deployment Engine.
   - **mNode Password**: The password for the Element QoSSIOC service. The QoSSIOC service default password is *solidfire*. If you have not created a custom password, you can create one from the Registration UI (`https://[management node IP]:9443`).
   - **vCenter User ID**: The user name for the vCenter admin with full Administrator role privileges.
   - **vCenter Password**: The password for the vCenter admin with full Administrator role privileges.
5. Click **OK**.

   The **mNode Status** field displays `UP` when the plug-in can successfully communicate with the service.

   Note: After you have configured a valid management node, these settings become the default. The mNode settings revert to the last known valid mNode settings until you provide settings for another valid management node. You must clear the settings for the configured management node before setting the credentials for a new management node.

Clearing management node settings for QoSSIOC

You can clear the configuration details of the Element management node (mNode) using the NetApp Element Configuration extension point. You must clear the settings for the configured management

node before configuring the credentials for a new management node. Clearing the mNode settings removes active QoSSIOC from the vCenter, cluster, and datastores.

Steps

1. Select **NetApp Element Configuration > mNode Settings**.
2. Click **Actions**.
3. In the resulting menu, select **Clear**.
4. Confirm the action.

   The **mNode Status** field displays `Not Configured` after the process is complete.

Related tasks

Changing the QoSSIOC service password on page 60

Changing the QoSSIOC service password

You can change the password for the QoSSIOC service using the registration user interface for the management node.

Before you begin

- You have cleared the configuration details of the Element management node (mNode) using the NetApp Element Configuration extension point.

  Note: You must clear the settings for the configured management node before configuring the credentials for a new management node. Clearing the mNode settings removes active QoSSIOC from the vCenter, cluster, and datastores.

- Your management node is powered on.

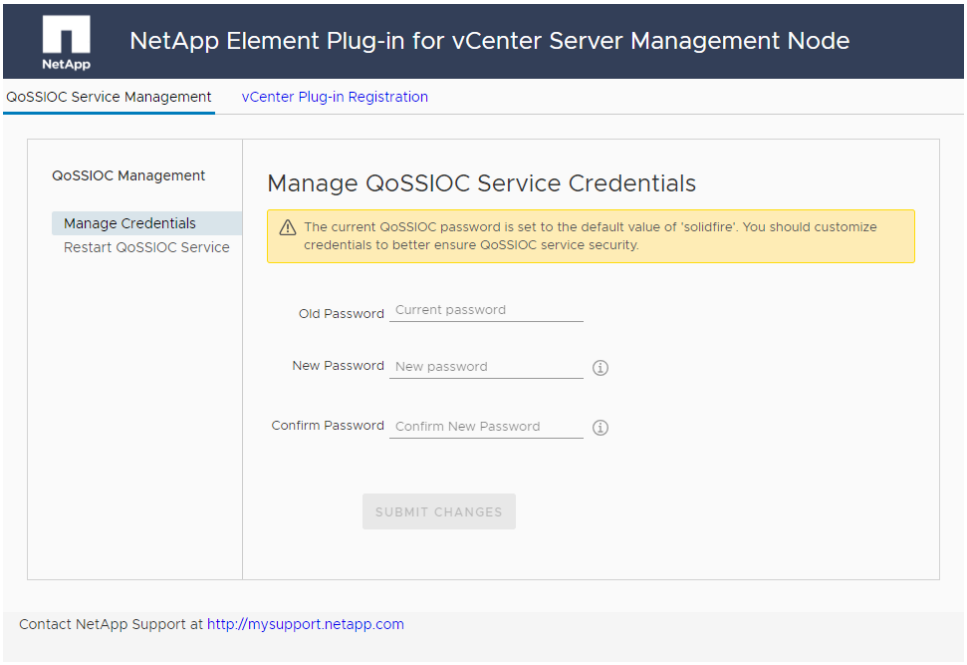
About this task

This process describes how to change the QoSSIOC password only. If you want to change the QoSSIOC user name, you can do so from the **mNode Settings** page of the NetApp Element Configuration extension point.

Steps

1. Enter the IP address for your management node in a browser, including the TCP port for registration: `https://[management node IP]:9443`.

   The registration UI displays the **Manage QoSSIOC Service Credentials** page for the plug-in.



2. Enter the following information:
   - **Old Password**: The current password of the QoSSIOC service. If you have not yet assigned a password, type the default password:

     `solidfire`
   - **New Password**: The new password for the QoSSIOC service.
   - **Confirm Password**: Enter the new password again.
3. Click **Submit Changes**.

   Note: The QoSSIOC service automatically restarts after you submit changes.
4. In your vSphere Web Client, select **NetApp Element Configuration > mNode Settings**.
5. Click **Actions**.
6. In the resulting menu, select **Configure**.
7. In the **Configure mNode Settings** dialog, enter the new password in the **mNode Password** field.
8. Click **OK**.

   The **mNode Status** field displays `UP` when the plug-in can successfully communicate with the service.

Related tasks

Clearing management node settings for QoSSIOC on page 59
Editing management node settings for QoSSIOC on page 59

Viewing QoSSIOC events

You can view QoSSIOC events from the NetApp Element Configuration extension point. A QoSSIOC event is reported when a VM that has a datastore with QoS enabled is reconfigured or issued a power or guest event.

Before you begin

- At least one cluster must be added and running.
- The QoSSIOC service must be configured and running using the **mNode Settings** page for the plug-in.
- At least one datastore must have QoSSIOC automation enabled.

About this task

QoSSIOC events are displayed from locally added clusters. In a Linked Mode environment, log into the vSphere Web Client that has the cluster added locally to view QoSSIOC events for that cluster.

Step

1. Select **NetApp Element Configuration > QoSSIOC Events**.

   The **QoSSIOC Events** page displays a list of events.

Related tasks

Configuring management node settings for QoSSIOC on page 58
Enabling QoSSIOC automation on page 90

QoSSIOC event details

You can view information about QoSSIOC events for each cluster on the **QoSSIOC Events** page of the **NetApp Element Configuration** extension point.

Date
The date and time of the QoSSIOC event.

Datastore Name
The user-defined datastore name.

Cluster IP
The IP address of the cluster containing the datastore from which the event originated.

Volume ID
The system-generated ID for the associated volume.

Min IOPs
The current minimum IOPS QoS setting of the volume.

Max IOPs
The current maximum IOPS QoS setting of the volume.

Burst IOPs
The current maximum burst QoS setting of the volume.

Burst Time

The length of time a burst is allowed.

Plug-in product information

You can find general information about the NetApp Element vCenter Plug-in version, build number, and IP address from the NetApp Element Configuration extension point. The **About** page includes options to download a service bundle and hide or restore the Getting Started pages that are available from both extension points.

Reporting

The **Reporting** tab gives you information about the cluster's components and provides an overview of how the cluster is performing. The **Reporting** page opens up into an overview of the cluster components and resources.

Reporting overview

You can view high-level cluster information for the selected cluster, including overall capacity, efficiency, and performance, on the **Overview** page of the **Reporting** tab from the **NetApp Element Management** extension point.

Cluster Capacity

The capacity remaining for block storage, metadata, and provisioned space. Move the pointer over the progress bar to see threshold information.

Cluster Information

Information specific to the cluster, such as cluster name, the version of NetApp Element software running on the cluster, MVIP and SVIP addresses, and the number of nodes, 4k IOPS, volumes, and sessions on the cluster.

- **Cluster Name**: The name for the cluster.
- **Storage IP (SVIP)**: The storage virtual IP address (SVIP).
- **Management IP (MVIP)**: The management virtual IP address (MVIP).
- **SVIP VLAN Tag**: The VLAN identifier for the master SVIP address.
- **MVIP VLAN Tag**: The VLAN identifier for the master MVIP address.
- **Node Count**: The number of active nodes in the cluster.
- **Cluster 4K IOPS**: The number of 4096 (4K) blocks that can be read/written by the cluster in a second.
- **Element OS Version**: The version of the NetApp Element software that the cluster is running.
- **Volume Count**: The total number of volumes, excluding virtual volumes, on the cluster.
- **Virtual Volume Count**: The total number of virtual volumes on the cluster.
- **iSCSI Sessions**: The iSCSI sessions that are connected to the cluster.
- **Fibre Channel Sessions**: The Fibre Channel sessions that are connected to the cluster.

Cluster Efficiency

Overall system capacity that is being utilized that takes into account thin provisioning, deduplication, and compression. The calculated benefit achieved on the cluster is calculated by comparing what the capacity utilization would be without thin provisioning, deduplication, and compression on a traditional storage device.

Protection Domains

A summary of protection domains monitoring for the cluster.

- **Selected Monitoring Level**: The protection domain resiliency levels as selected by the user. Possible values are **Chassis** or **Node**. Green indicates that the cluster is capable of the selected monitoring level. Red indicates that the cluster is no longer capable of the selected monitoring level and corrective action is needed.
- **Remaining Block Capacity**: Indicates the remaining block capacity that can be used while maintaining the selected resiliency level.
- **Metadata Capacity**: Indicates if there is sufficient metadata capacity to heal from failure while also maintaining undisrupted data availability. **Normal** (green) indicates that the cluster has sufficient metadata to maintain the selected monitoring level. **Full** (red) indicates that the cluster is no longer capable of the selected monitoring level and corrective action is needed.

Provisioned IOPS

A summary of how volume IOPS might be overprovisioned on the cluster. Provisioned IOPS calculations are determined by the sum of the total minimum IOPS, maximum IOPS, and burst IOPS for all volumes on the cluster divided by the maximum IOPS rated for the cluster.

Note: For example, if there are four volumes in the cluster, each with minimum IOPS of 500, maximum IOPS of 15,000, and burst IOPS of 15,000, the total number of minimum IOPS would be 2,000, total maximum IOPS would be 60,000, and total burst IOPS would be 60,000. If the cluster is rated at maximum IOPS of 50,000,then the calculations would be the following:

- Minimum IOPS: 2000/50000 = 0.04x
- Maximum IOPS: 60000/50000 = 1.20x
- Burst IOPS: 60000/50000 = 1.20x 1.00x

1.00x is the baseline at which provisioned IOPS is equal to the rated IOPS for the cluster.

Cluster Health

The hardware, capacity, and security components of the health of the cluster. Color codes indicate the following:

- Green: Healthy
- Yellow: Critical
- Red: Error

Cluster Input/Output

The I/O currently running on the cluster. The values are calculated based on the previous I/O measurement against the current I/O measurements. These are the measurements shown in the graph:

- **Total**: The combined read and write IOPS occurring in the system.
- **Read**: The number of read IOPS occurring.
- **Write**: The number of write IOPS.

Cluster Throughput

The bandwidth activity for read, write, and total bandwidth on the cluster:

- **Total**: The total MB/s used for both read and write activity in the cluster.
- **Read**: The read activity in MB/s for the cluster.

- **Write**: The write activity in MB/s for the cluster.

Performance Utilization

The percentage of cluster IOPS being consumed. For example, a 250K IOPS cluster running at 100K IOPS would show 40% consumption.

Viewing event logs

You can review event logs for operations performed on the selected cluster along with cluster faults that might occur. Most errors are resolved automatically by the system. Other faults might require manual intervention.

Steps

1. Select **NetApp Element Management > Reporting**.

   Note: If two or more clusters are added, the cluster you intend to use for the task must be selected.

2. Click **Event Log**.

   The page displays a list of all events on the cluster.

3. Select an individual event you want to review.

4. Select **Details**.

   The message displays the cluster event details.

Event log

You can view information about events detected in the system on the **Event Log** page of the **Reporting** tab from the **NetApp Element Management** extension point.

Event 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.

Service ID

The ID of the service that reported the event (if applicable). The value is `0` (zero) if the fault is not associated with a service.

Node

The ID of the node that reported the event (if applicable).

Drive ID

The ID of the drive that reported the event (if applicable).

Event Time

The date and time the event occurred.

Event types

The system reports multiple types of events; each event is an operation that the system has completed. Events can be routine, normal events or events that require administrator attention. The **Event Type** 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 might 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.

installEvent

Automatic software installation events. Software is being automatically installed on a pending node.

iSCSIEvent

Events related to iSCSI issues in the system.

limitEvent

Events related to the number of volumes or virtual volumes in an account or in the cluster nearing the maximum allowed.

networkEvent

Events related to the status of virtual networking.

platformHardwareEvent

Events related to issues detected on hardware devices.

remoteClusterEvent

Events related to remote cluster pairing.

schedulerEvent

Events related to scheduled snapshots.

serviceEvent

Events related to system service status.

sliceEvent

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

snmpTrapEvent

Events related to SNMP traps.

statEvent

Events related to system statistics.

tsEvent

Events related to the system transport service.

unexpectedException

Events related to unexpected system exceptions.

vasaProviderEvent

Events related to a VASA (vSphere APIs for Storage Awareness) Provider.

Alerts

Alerts are cluster faults or errors and are reported as they occur. Alerts can be information, warnings, or errors and are a good indicator of how well the cluster is running. Most errors resolve themselves automatically; however, some might require manual intervention.

You can view information about individual system alerts on the **Alerts** page of the **Reporting** tab from the **NetApp Element Management** extension point.

After the system resolves an alert, all information about the alert including the date it was resolved is moved to the **Resolved** view.

The following list describes the columns on the page.

ID

Unique ID for a cluster alert.

Severity

- `warning`: A minor issue that might soon require attention. System upgrades are still allowed at this severity level.
- `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

- `node`: Fault affecting an entire node.
- `drive`: Fault affecting an individual drive.
- `cluster`: Fault affecting the entire cluster.
- `service`: Fault affecting a service on the cluster.
- `volume`: Fault affecting a volume on the cluster.

Node

Node ID for the node that this fault refers to. Included for node and drive faults, otherwise set to - (dash).

Drive ID

Drive ID for the drive that this fault refers to. Included for drive faults, otherwise set to - (dash).

Error Code

A descriptive code that indicates what caused the fault.

Details

Detailed description of the fault.

Time

This heading is only visible in **Active filter** view. The date and time the fault was logged.

Resolution Date

This heading is only visible in **Resolved filter** view. The date and time the fault was resolved.

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 why the alert was generated.

The following list outlines the different types of system alerts.

availableVirtualNetworkIPAddressesLow

The number of virtual network addresses in the block of IP addresses is low. To resolve this fault, add more IP addresses to the block of virtual network addresses.

blockClusterFull

There is not enough free block storage space to support a single node loss. To resolve this fault, add another storage node to the storage cluster.

blockServiceTooFull

A block service is using too much space. To resolve this fault, add more provisioned capacity.

blockServiceUnhealthy

A block service has been detected as unhealthy. The system is automatically moving affected data to other healthy drives.

clusterCannotSync

There is an out-of-space condition and data on the offline block storage drives cannot be synced to drives that are still active. To resolved this fault, add more storage.

clusterFull

There is no more free storage space in the storage cluster. To resolve this fault, add more storage.

clusterIOPSAreOverProvisioned

Cluster IOPS are over provisioned. The sum of all minimum QoS IOPS is greater than the expected IOPS of the cluster. Minimum QoS cannot be maintained for all volumes simultaneously.

disconnectedClusterPair

A cluster pair is disconnected or configured incorrectly.

disconnectedRemoteNode

A remote node is either disconnected or configured incorrectly.

disconnectedSnapMirrorEndpoint

A remote SnapMirror endpoint is disconnected or configured incorrectly.

driveAvailable

One or more drives are available in the cluster. In general, all clusters should have all drives added and none in the available state. If this fault appears unexpectedly, contact NetApp Support. To resolve this fault, add any available drives to the storage cluster.

driveFailed

One or more drives have failed. Contact NetApp Support to have the drive replaced.

driveWearFault

A drive's remaining life has dropped below thresholds, but it is still functioning. To resolve this fault, replace the drive soon.

duplicateClusterMasterCandidates

More than one storage cluster master candidate has been detected. Contact NetApp Support for assistance.

ensembleDegraded

Network connectivity or power has been lost to one or more of the ensemble nodes. To resolve this fault, restore network connectivity or power.

exception

A fault reported that is other than a routine fault. These faults are not automatically cleared from the fault queue. Contact NetApp Support for assistance.

failedSpaceTooFull

A block service is not responding to data write requests. This causes the slice service to run out of space to store failed writes. To resolve this fault, restore block services functionality to allow writes to continue normally and failed space to be flushed from the slice service.

fanSensor

A fan sensor has failed or is missing. Contact NetApp Support for assistance.

fibreChannelAccessDegraded

A Fibre Channel node is not responding to other nodes in the storage cluster over its storage IP for a period of time. In this state, the node will then be considered unresponsive and generate a cluster fault.

fibreChannelAccessUnavailable

All Fibre Channel nodes are unresponsive. The node IDs are displayed.

fibreChannelConfig

This cluster fault indicates one of the following conditions:

- There is an unexpected Fibre Channel port on a PCI slot.
- There is an unexpected Fibre Channel HBA model.
- There is a problem with the firmware of a Fibre Channel HBA.
- A Fibre Channel port is not online.
- There is a persistent issue configuring Fibre Channel passthrough.

Contact NetApp Support for assistance.

fileSystemCapacityLow

There is insufficient space on one of the filesystems.

To resolve this fault, add more capacity to the filesystem.

FIPS drives mismatched

A non-FIPS drive has been physically inserted into a FIPS capable storage node or a FIPS drive has been physically inserted into a non-FIPS storage node. A single fault is generated per node and lists all drives affected.

To resolve this fault, remove or replace the drive or drives in question.

FIPS drives out of compliance

The system has detected that Encryption at Rest was disabled after the FIPS Drives feature was enabled. This fault is also generated when the FIPS Drives feature is enabled and a non-FIPS drive or node is present in the storage cluster.

To resolve this fault, enable Encryption at Rest or remove the non-FIPS hardware from the storage cluster.

fipsSelfTestFailure

The FIPS subsystem has detected a failure during the self test.

Contact NetApp Support for assistance.

hardwareConfigMismatch

This cluster fault indicates one of the following conditions:

- The configuration does not match the node definition.
- There is an incorrect drive size for this type of node.
- An unsupported drive has been detected.
- There is a drive firmware mismatch.
- The drive encryption capable state does not match the node.

Contact NetApp Support for assistance.

inconsistentBondModes

The bond modes on the VLAN device are missing. This fault will display the expected bond mode and the bond mode currently in use.

inconsistentInterfaceConfiguration

The interface configuration is inconsistent.

To resolve this fault, ensure the node interfaces in the storage cluster are consistently configured.

inconsistentMtus

This cluster fault indicates one of the following conditions:

- Bond1G mismatch: Inconsistent MTUs have been detected on Bond1G interfaces.
- Bond10G mismatch: Inconsistent MTUs have been detected on Bond10G interfaces.

This fault displays the node or nodes in question along with the associated MTU value.

inconsistentRoutingRules

The routing rules for this interface are inconsistent.

inconsistentSubnetMasks

The network mask on the VLAN device does not match the internally recorded network mask for the VLAN. This fault displays the expected network mask and the network mask currently in use.

incorrectBondPortCount

The number of bond ports is incorrect.

invalidConfiguredFibreChannelNodeCount

One of the two expected Fibre Channel node connections is degraded. This fault appears when only one Fibre Channel node is connected.

irqBalanceFailed

An exception occurred while attempting to balance interrupts.

Contact NetApp Support for assistance.

memoryUsageThreshold

Memory usage is above normal.

Contact NetApp Support for assistance.

metadataClusterFull

There is not enough free metadata space to support a single node loss.

To resolve this fault, add another storage node to the storage cluster.

mtuCheckFailure

A network device is not configured for the proper MTU size.

To resolve this fault, ensure that all network interfaces and switch ports are configured for jumbo frames (MTUs up to 9000 bytes in size).

networkConfig

This cluster fault indicates one of the following conditions:

- An expected interface is not present.
- A duplicate interface is present.
- A configured interface is down.

- A network restart is required.

Contact NetApp Support for assistance.

networkErrorsExceedThreshold

This cluster fault indicates one of the following conditions:

- The number of frame errors is above normal.
- The number of CRC errors is above normal.

To resolve this fault, replace the network cable connected to the interface reporting these errors.

Contact NetApp Support for assistance.

noAvailableVirtualNetworkIPAddresses

There are no available virtual network addresses in the block of IP addresses. No more storage nodes can be added to the cluster.

To resolve this fault, add more IP addresses to the block of virtual network addresses.

nodeOffline

Element software cannot communicate with the specified node.

notUsingLACPBondMode

LACP bonding mode is not configured.

To resolve this fault, use LACP bonding when deploying storage nodes; clients might experience performance issues if LACP is not enabled and properly configured.

ntpServerUnreachable

The storage cluster cannot communicate with the specified NTP server or servers.

To resolve this fault, check the configuration for the NTP server, network, and firewall.

ntpTimeNotInSync

The difference between storage cluster time and the specified NTP server time is too large. The storage cluster cannot correct the difference automatically. To resolve this fault, use NTP servers that are internal to your network, rather than the installation defaults. If you are using internal NTP servers and the issue persists, contact NetApp Support for assistance.

nvramDeviceStatus

An NVRAM device has an error, is failing, or has failed.

Contact NetApp Support for assistance.

powerSupplyError

This cluster fault indicates one of the following conditions:

- A power supply is not present.
- A power supply has failed.
- A power supply input is missing or out of range.

To resolve this fault, verify that redundant power is supplied to all nodes. Contact NetApp Support for assistance.

provisionedSpaceTooFull

The overall provisioned capacity of the cluster is too full.

To resolve this fault, add more provisioned space, or delete and purge volumes.

remoteRepAsyncDelayExceeded

The configured asynchronous delay for replication has been exceeded.

remoteRepClusterFull

The volumes have paused remote replication because the target storage cluster is too full.

remoteRepSnapshotClusterFull

The volumes have paused remote replication of snapshots because the target storage cluster is too full.

To resolve this fault, free up some space on the target storage cluster.

remoteRepSnapshotsExceededLimit

The volumes have paused remote replication of snapshots because the target storage cluster volume has exceeded its snapshot limit.

scheduleActionError

One or more of the scheduled activities ran, but failed.

The fault clears if the scheduled activity runs again and succeeds, if the scheduled activity is deleted, or if the activity is paused and resumed.

sensorReadingFailed

The Baseboard Management Controller (BMC) self-test failed or a sensor could not communicate with the BMC.

Contact NetApp Support for assistance.

serviceNotRunning

A required service is not running.

Contact NetApp Support for assistance.

sliceServiceTooFull

A slice service has too little provisioned capacity assigned to it.

To resolve this fault, add more provisioned capacity.

sliceServiceUnhealthy

The system has detected that a slice service is unhealthy and is automatically decommissioning it.

sshEnabled

The SSH service is enabled on one or more nodes in the storage cluster.

To resolve this fault, disable the SSH service on the appropriate node or nodes or contact NetApp Support for assistance.

sslCertificateExpiration

The SSL certificate associated with this node has expired.

To resolve this fault, renew the SSL certificate. If needed, contact NetApp Support for assistance.

tempSensor

A temperature sensor is reporting higher than normal temperatures. This fault can be triggered in conjunction with powerSupplyError or fanSensor faults.

To resolve this fault, check for airflow obstructions near the storage cluster. If needed, contact NetApp Support for assistance.

upgrade

An upgrade has been in progress for more than 24 hours.

to resolve this fault, resume the upgrade or contact NetApp Support for assistance.

unbalancedMixedNodes

A single node accounts for more than one-third of the storage cluster's capacity.

Contact NetApp Support for assistance.

unresponsiveService

A service has become unresponsive.

Contact NetApp Support for assistance.

virtualNetworkConfig

This cluster fault indicates one of the following conditions:

- An interface is not present.
- There is an incorrect namespace on an interface.
- There is an incorrect netmask.
- There is an incorrect IP address.
- An interface is not up and running.
- There is a superfluous interface on a node.

Contact NetApp Support for assistance.

volumeDegraded

Secondary volumes have not finished replicating and synchronizing. The message is cleared when the synchronizing is complete.

volumesOffline

One or more volumes in the storage cluster are offline.

Contact NetApp Support for assistance.

Related information

NetApp Support

iSCSI sessions

You can view information about iSCSI sessions that are connected to the selected cluster on the **iSCSI Sessions** page of the **Reporting** tab from the **NetApp Element Management** extension point.

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 (-) will be 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 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.

Running tasks

You can view the progress and completion status of running tasks that are being reported by `ListSyncJobs` and `ListBulkVolumeJobs` API methods. You can view running tasks on the **Running Tasks** page of the **Reporting** tab from the **NetApp Element Management** extension point.

Task Type

The type of sync job or bulk volume job.

Possible values:

- `block`
- `clone`
- `read`
- `remote`
- `slice`
- `write`

Node

Specifies the ID of the node onto which the clone is being written. This ID is only present if the task type is `clone`.

Description

List of objects describing sync processes or an array of information for each bulk volume job currently running in the system.

Current Progress

Number of bytes the clone has processed in the source volume. This information is only present if the task type is `clone` or `slice`.

Elapsed Time

The time elapsed, in seconds, since the job started.

Remaining Time

Estimated time, in seconds, to complete the operation.

Volume performance

You can view performance information for all volumes in the selected cluster from the **Volume Performance** page on the **Reporting** tab of the **NetApp Element Management** extension point.

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.

ID

The system-generated ID for the volume.

Name

The name given to the volume when it was created.

Account

The name of the account assigned to the volume.

Access Groups

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

Volume Utilization %

A percentage value that describes how much the client is using the volume.

Possible values:

- 0 = Client is not using the volume
- 100 = Client is using the max
- >100 = Client is using the burst

Total IOPS

The total number of IOPS (read and write) currently being executed against the volume.

Read IOPS

The total number of read IOPS currently being executed against the volume.

Write IOPS

The total number of write IOPS currently being executed against the volume.

Total Throughput

The total amount of throughput (read and write) currently being executed against the volume.

Read Throughput

The total amount of read throughput currently being executed against the volume.

Write Throughput

The total amount of write throughput currently being executed against the volume.

Total Latency (ms)

The average time, in microseconds, to complete read and write operations to a volume.

Read Latency (ms)

The average time, in microseconds, to complete read operations to the volume in the last 500 milliseconds.

Write Latency (ms)

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.

Management

The **Management** tab enables you to create and manage datastores, volumes, accounts, access groups, initiators, and volume QoS policies.

Datastore management

The NetApp Element Plug-in for vCenter Server enables you to manage datastores that are created on Element volumes. You can create, extend, clone, share, or delete datastores. You can also use VAAI UNMAP to allow a cluster to reclaim freed block space from thinly provisioned VMFS datastores.

Datastore operations can be monitored for completion using task monitoring in vSphere.

Note: To create and manage datastores, you must first create at least one user account in **Management > Accounts**. If you want to use QoSSIOC service with datastores, you must first configure settings on the **mNode Settings** page from the NetApp Element Configuration extension point.

Related tasks

Creating an account on page 109
Configuring management node settings for QoSSIOC on page 58

Creating a datastore

You can create a datastore using the **NetApp Element Management** extension point.

Before you begin

- At least one cluster must be added and running.

  Note: If you are using vCenter Linked Mode, be sure that you have added your cluster with the correct vCenter Server.

- At least one user account must be created.
- At least one host must be connected to the vCenter Server.

About this task

Because datastores are created using the highest VMFS version supported by the selected ESXi host, all cluster members should run the same version of vSphere and ESXi to avoid VMFS compatibility issues.

Note: If you want to use QoSSIOC service with datastores, you must first configure settings on the **mNode Settings** page from the **NetApp Element Configuration** extension point.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Datastore** page, click **Create Datastore**.
3. In the **Create Datastore** dialog box, enter a name for the datastore.

Tip: Use a unique name for each datastore in a data center. For multiple cluster or vCenter Server environments, use descriptive naming best practices.

4. Click **Next**.
5. Select one or more required hosts for the datastore.

   Note: You need at least one connected host before you can create a new datastore. If your host has multiple initiators, select an initiator or the host to select all initiators. If you are using vCenter Linked Mode, only hosts available to the vCenter Server to which the cluster is assigned are available to select.

6. Click **Next**.
7. In the **Configure Volume** pane, select an existing volume and proceed to the next step, or create a new volume for the new datastore:

   a. Enter a name for the volume that backs the datastore.

   b. Select a user account from the account drop-down list.

      Note: You need at least one existing user account before you can create a new volume.

   c. Enter the total size of the volume you want to create.

      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

      Note: By default, 512 byte emulation is set to ON for all the new volumes.

   d. In the **Quality of Service** area, do one of the following:

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

        Attention: After you enable datastore QoSSIOC settings, any QoS settings at the volume level are overridden.

        Attention: 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 **Next**.
9. Configure authorization type for host access by selecting one of the following:

   - **Use Volume Access Group**
   - **Use CHAP**

   Note: Use the volume access group authorization type to explicitly limit which initiators can see volumes. Use CHAP for secure secret-based access with no limits on initiators.

10. Click **Next**.
11. If you selected **Use Volume Access Group**, configure the volume access groups for the selected hosts.

The volume access groups listed in **Required by Selected Initiators** are already associated with one or more of the host initiators you selected in an earlier step.

a. Select additional volume access groups or create new ones to associate with available initiators:

- **Available**: Other volume access group options in the cluster.
- **Create New Access Group**: Enter the name of the new access group and click **Add**.

b. Click **Next**.

c. In the **Configure Hosts' Access** pane, associate available host initiators (IQN or WWPN) with the volume access groups you selected in the previous pane. If a host initiator is already associated with a volume access group, the field is read-only for that initiator. If a host initiator does not have a volume access group association, select an option from the drop-down list next to the initiator.

d. Click **Next**.

12. If you want to enable QoSSIOC automation, click the **Enable QoS & SIOC** check box to select it and then configure the QoSSIOC settings.

    Note: If the QoSSIOC service is not available, you must first configure settings on the **mNode Settings** page from the **NetApp Element Configuration** extension point.

    a. Select **Enable QoS & SIOC**.

    b. Configure the **Burst Factor**.

    Note: The burst factor is a multiple of the IOPS limit (SIOC) setting for the VMDK. If you change the default, make sure to use a burst factor value that will not exceed the maximum burst limit for an Element volume when the burst factor value is multiplied by the IOPS limit for any VMDK.

    c. Optional: Select **Override Default QoS** and configure the settings.

    Note: If the **Override Default QoS** setting is disabled for the datastore, the `Shares` and `Limit IOPS` values are automatically set based on the default SIOC settings of each VM.

    Tip: Do not customize the SIOC share limit without also customizing the SIOC IOPS limit.

13. Click **Next**.

14. Confirm the selections and click **Finish**.

15. You can use task monitoring in vSphere to view the progress of the task. Refresh the view if the datastore does not appear in the list.

Related concepts

Using object naming best practices when managing multiple clusters on page 49

Related tasks

Adding a cluster on page 50
Creating an account on page 109
Configuring management node settings for QoSSIOC on page 58

Viewing the datastore list

You can view available datastores on the **Datastores** page from the **NetApp Element Management** extension point.

Steps

1. Select **NetApp Element Management** > **Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. The **Datastores** page appears and shows all current datastores on NetApp Element and ESXi.

   Note: Datastores spanning multiple volumes (mixed datastores) are not listed. Datastore views will only show datastores that are available on ESXi hosts from the selected NetApp Element cluster.

Datastore details

You can view the information for all datastores on the cluster on the **Datastores** page of the **Management** tab from the **NetApp Element Management** extension point.

Name

The name assigned to the datastore.

Host Name(s)

The address of each associated host device.

Status

The possible values `Accessible` or `Inaccessible` indicate whether or not the datastore is currently connected to vSphere.

Type

The VMware file system datastore type.

Volume Name

The name assigned to the associated volume.

Volume NAA

Globally unique SCSI device identifier for the associated volume in NAA IEEE Registered Extended format.

Total Capacity (GB)

Total formatted capacity of the datastore.

Free Capacity (GB)

Space that is available for the datastore.

QoSSIOC Automation

Indicates whether or not QoSSIOC automation is enabled. Possible values:

- `Enabled`: QoSSIOC is enabled.
- `Disabled`: QoSSIOC is not enabled.
- `Max Exceeded`: Volume Max QoS has exceeded the limit value specified.

Extending a datastore

You can extend a datastore to increase volume size using the NetApp Element Management extension point. Extending the datastore also extends the VMFS volume related to that datastore.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Datastores** page, select the check box for the datastore you want to extend.
3. Click **Actions**.
4. In the resulting menu, click **Extend**.
5. In the **New Datastore Size** field, type the required size for the new datastore and select GB or GiB.

   Note: Extending the datastore will consume the entire volume's size. The new datastore size cannot exceed the unprovisioned space available on the selected cluster or the maximum volume size the cluster allows.

6. Click **OK**.
7. Refresh the page if needed until the updated datastore appears in the list.

Cloning a datastore

The NetApp Element Plug-in for vCenter Server provides the functionality to clone datastores, which includes mounting the new datastore to the desired ESXi server or cluster. You can name the datastore clone and configure its QoSSIOC, volume, host, and authorization type settings.

Before you begin

- At least one cluster must be added and running.

  Note: If you are using vCenter Linked Mode, be sure that you have added your cluster with the correct vCenter Server.

- Available unprovisioned space must be equal to or more than the source volume size.
- At least one host must be connected to vCenter Server.
- At least one user account must be created.

About this task

If virtual machines exist on the source datastore, virtual machines on the clone datastore will be brought into the inventory with new names.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Datastores** page, select the check box for the datastore you want to clone.

3. Click **Actions**.
4. In the resulting menu, click **Clone**.

   Note: If you attempt to clone a datastore that contains virtual machines with attached disks not located on the selected datastore, copies of the virtual machines on the cloned datastore will not be added to the virtual machine inventory.

5. Enter a datastore name.

   Tip: Use a unique name for each datastore in a data center. For multiple cluster or vCenter Server environments, use descriptive naming best practices.

6. Click **Next**.
7. Select one or more required hosts for the datastore.

   Note: You need at least one connected host before you can create a new datastore. If your host has multiple initiators, select an initiator or the host to select all initiators. If you are using vCenter Linked Mode, only hosts available to the vCenter Server to which the cluster is assigned are available to select.

8. Click **Next**.
9. In the **Configure Volume** pane, do the following:

   Note: Volume size for the clone datastore matches the size of the volume backing the source datastore. By default, 512 byte emulation is set to ON for all the new volumes.

   a. Enter a name for the new NetApp Element volume that backs the clone datastore.

   b. Select a user account from the account drop-down list.

      Note: You need at least one existing user account before you can create a new volume.

   c. In the **Quality of Service** area, do one of the following:

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

        Attention:

        - After you enable datastore QoSSIOC settings, any QoS settings at the volume level are overridden.
        - 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.

10. Click **Next**.
11. Configure authorization type for host access by selecting one of the following options:

    - **Use Volume Access Group**
    - **Use CHAP**

    Note: Use the volume access group authorization type to explicitly limit which initiators can see volumes. Use CHAP for secure secret-based access with no limits on initiators.

12. Click **Next**.

13. If you selected **Use Volume Access Group**, configure the volume access groups for the selected hosts.

    The volume access groups listed in **Required by Selected Initiators** are already associated with one or more of the host initiators you selected in an earlier step.

    a. Select additional volume access groups or create new ones to associate with available initiators:

       - **Available**: Other volume access group options in the cluster.
       - **Create New Access Group**: Enter the name of the new access group and click **Add**.

    b. Click **Next**.

    c. In the **Configure Hosts' Access** pane, associate available host initiators (IQN or WWPN) with the volume access groups you selected in the previous pane.

       If a host initiator is already associated with a volume access group, the field is read-only for that initiator. If a host initiator does not have a volume access group association, select an option from the drop-down list next to the initiator.

    d. Click **Next**.

14. If you want to enable QoSSIOC automation, click the **Enable QoS & SIOC** check box to select it and then configure the QoSSIOC settings.

    Note: If the QoSSIOC service is not available, you must first configure settings in the **mNode Settings** page in the **NetApp Element Configuration** extension point.

    a. Select **Enable QoS & SIOC**.

    b. Configure the **Burst Factor**.

       Note: The burst factor is a multiple of the IOPS limit (SIOC) setting for the VMDK. If you change the default, make sure to use a burst factor value that will not exceed the maximum burst limit for a NetApp Element volume when the burst factor value is multiplied by the IOPS limit for any VMDK.

    c. Optional: Select **Override Default QoS** and configure the settings.

       Note: If the **Override Default QoS** setting is disabled for the datastore, the `Shares` and `Limit IOPS` values are automatically set based on the default SIOC settings of each VM.

       Tip: Do not customize the SIOC share limit without also customizing the SIOC IOPS limit.

15. Click **Next**.

16. Confirm the selections and click **Finish**.

17. Refresh the view if the datastore clone does not appear in the list.

Related concepts

Using object naming best practices when managing multiple clusters on page 49

Related tasks

Adding a cluster on page 50
Creating an account on page 109

Sharing a datastore

You can share a datastore with one or more hosts using the NetApp Element Management extension point.

Before you begin

- At least one cluster must be added and running.

  Note: If you are using vCenter Linked Mode, be sure that you have added your cluster with the correct vCenter Server.

- There must be more than one host under the selected data center.

  Note: Datastores can be shared only among hosts within the same data center.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Datastores** page, select the check box for the datastore you want to share.
3. Click **Actions**.
4. In the resulting menu, click **Share**.
5. Configure authorization type for host access by selecting one of the following:

   - **Use Volume Access Group**
   - **Use CHAP**

   Note: Use the volume access group authorization type to explicitly limit which initiators can see volumes. Use CHAP for secure secret-based access with no limits on initiators.

6. Click **Next**.
7. Select one or more required hosts for the datastore.

   Note: You need at least one connected host before you can create a new datastore. If your host has multiple initiators, select an initiator or all initiators by selecting the host. If you are using vCenter Linked Mode, only hosts available to the vCenter Server to which the cluster is assigned are available to select.

8. Click **Next**.
9. If you selected **Use Volume Access Group**, configure the volume access groups for the selected hosts. The volume access groups listed in **Required by Selected Initiators** are already associated with one or more of the host initiators you selected in an earlier step.

   a. Select additional volume access groups or create new ones to associate with available initiators:

      - **Available**: Other volume access group options in the cluster.
      - **Create New Access Group**: Enter the name of the new access group and click **Add**.

   b. Click **Next**.

c. In the **Configure Hosts' Access** pane, associate available host initiators (IQN or WWPN) with the volume access groups you selected in the previous pane.

   If a host initiator is already associated with a volume access group, the field is read-only for that initiator. If a host initiator does not have a volume access group association, select an option from the drop-down list next to the initiator.

10. Confirm the selections and click **Finish**.
11. Refresh the list after the share datastore task is complete to verify hosts for the datastore.

Performing VAAI UNMAP

The VAAI UNMAP feature allows a cluster to reclaim freed block space from thinly provisioned VMFS5 datastores.

Before you begin

1. Ensure that the datastore you are using for the task is VMFS5 or earlier. VAAI UNMAP is unavailable for VMFS6 because ESXi performs the task automatically
2. Ensure that the ESXi host system settings are enabled for VAAI UNMAP:

   ```
   esxcli system settings advanced list -o/VMFS3/EnableBlockDelete
   ```

   The integer value must be set to 1 to enable.
3. If the ESXi host system settings are not enabled for VAAI UNMAP, set the integer value to 1 with this command:

   ```
   esxcli system settings advanced set -i 1 -o /VMFS3/EnableBlockDelete
   ```

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.
2. From the **Datastores** page, select the check box for the datastore on which you want to use VAAI UNMAP.
3. Click **Actions**.
4. In the resulting menu, click **VAAI Unmap**.
5. Select a host by name or IP address.
6. Enter the host user name and password.
7. Confirm the selections and click **OK**.

Related information

VMware VAAI blog article

Deleting a datastore

You can delete a datastore using the NetApp Element Management extension point. This operation permanently deletes all the files associated with the VMs on the datastore that you want to delete. The plug-in does not delete datastores that contain registered VMs.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Datastores** page, select the check box for the datastore you want to delete.
3. Click **Actions**.
4. In the resulting menu, click **Delete**.
5. Optional: If you want to delete the NetApp Element volume that is associated with the datastore, select the **Delete associated volume** check box.

   Note: You can also choose to retain the volume and later associate it with another datastore.

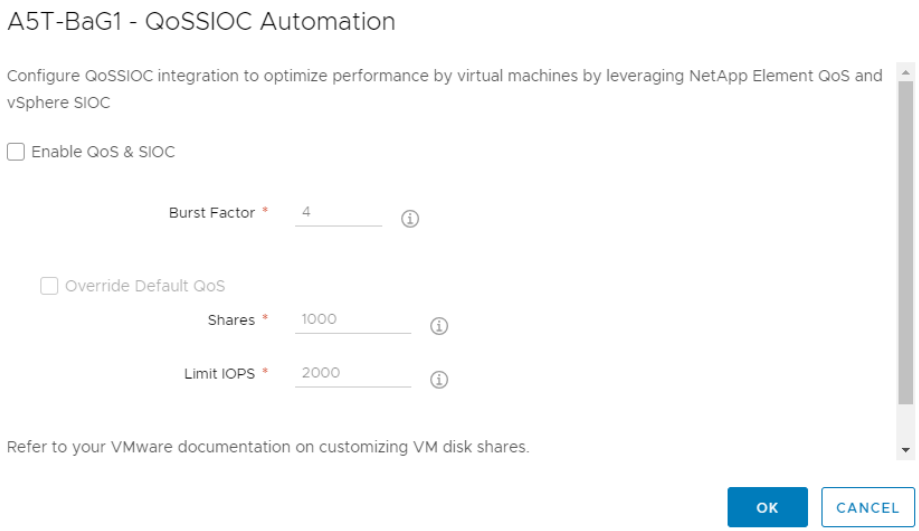
6. Click **Yes**.

QoSSIOC automation

The NetApp Element Plug-in for vCenter Server allows, as an optional setting, automatic quality of service (QoS) based on storage I/O control (SIOC) settings of all VMs on a datastore.

QoS and SIOC integration (QoSSIOC), which can be enabled for datastores in the user interface, runs a scan of all SIOC settings on all associated VMs. QoSSIOC automation is triggered by VM power, guest, and reconfiguration activity. The QoSSIOC service uses the sum of all SIOC reservations or shares and the sum of IOPS limits to determine minimum and maximum QoS for the underlying volume of each datastore. A configurable burst factor is also available.

The following dialog box appears during QoSSIOC configuration for a datastore:



Enable QoS & SIOC

Enables the automatic monitoring of SIOC values for each VMDK on a datastore and sets QoS values for the underlying volume according to those values.

Burst Factor

Multiplier of the sum of SIOC IOPS limit values from each VDMK that determines the burst IOPS contribution for the underlying volume.

Override Default QoS

Enables the use of Shares and Limit IOPS values. These values can be used when SIOC settings for each VM are set to default.

Shares

The contribution of minimum IOPS from each VDMK if the SIOC settings are set to default.

Limit IOPS

The contribution of maximum IOPS from each VDMK if the SIOC settings are set to default.

When SIOC settings for a VMDK are at the default shares level of Normal and the default IOPS limit of Unlimited, the Shares and Limit IOPS values contribute toward the total QoS for the underlying volume. If the SIOC settings for the VMDK are not at default levels, SIOC shares contribute to Min QoS and SIOC IOPS limit values contribute to Max QoS for the underlying volume.

Note: It is possible to set a reservation value through vSphere API. If a reservation value is set for a VMDK, shares are ignored and the reservation value is used instead.

Enabling QoSSIOC automation

You can enable QoSSIOC automation and customize virtual machine disk (VMDK) performance levels using the NetApp Element Management extension point.

Before you begin

You have configured the QoSSIOC service settings on the **mNode Settings** page from the **NetApp Element Configuration** extension point.

Note: If you are using QoS policies, do not enable QoSSIOC. QoSSIOC will override any volume QoS settings.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the status button in the **QoSSIOC Automation** column for the selected datastore.

   Tip: Ensure that the datastore does not have QoSSIOC integration enabled on another vCenter to prevent unexpected changes in QoS.

3. Select **Enable QoS & SIOC**.

4. Configure the **Burst Factor**.

   Note: The burst factor is a multiple of the IOPS limit (SIOC) setting for the VMDK. If you change the default, make sure to use a burst factor value that will not exceed the maximum

burst limit for a NetApp Element software-based volume when the burst factor value is multiplied by the IOPS limit for any VMDK.

5. Optional: Select **Override Default QoS** and configure the settings.

   Note: If the **Override Default QoS** setting is disabled for the datastore, the `Shares` and `Limit IOPS` values are automatically set based on the default SIOC settings of each VM.

   Tip: Do not customize the SIOC share limit without also customizing the SIOC IOPS limit.

6. Click **OK**.

   Note: When you enable the QoSSIOC Automation for a datastore, the button changes from `Disabled` to `Enabled`.

Related tasks

Configuring management node settings for QoSSIOC on page 58

Disabling QoSSIOC integration

You can disable QoSSIOC integration using the NetApp Element Management extension point.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the button in the **QoSSIOC Automation** column for the selected datastore.
3. Clear the **Enable QoS & SIOC** check box to disable the integration.

   Note: Clearing the **Enable QoS & SIOC** check box automatically disables the **Override Default QoS** option.

4. Click **OK**.

Volume management

Storage is provisioned in the NetApp Element system as volumes. Volumes are block devices accessed over the network using iSCSI or Fibre Channel clients.

The NetApp Element Plug-in for vCenter Server enables you to create, view, edit, delete, clone, backup or restore volumes for user accounts. You can also manage each volume on a cluster, and add or remove volumes in volume access groups.

Creating a volume

You can create a new 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. You can also specify QoS settings for a volume during creation.

Before you begin

- At least one cluster must be added and running.
- A user account has been created.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, click **Create Volume**.
4. Enter a name for the volume.

   Tip: Use descriptive naming best practices. This is especially important if multiple clusters or vCenter Servers are used in your environment.

5. Enter the total size of the volume you want to create.

   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

   Note: By default, 512 byte emulation is set to ON for all the new volumes. VMware requires 512e for disk resources. If 512e is not enabled, a VMFS cannot be created.

6. Select a user account from the **Account** drop-down list.
7. In the **Quality of Service** area, do one of the following:

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

   Attention:

   - After you enable datastore QoSSIOC settings, any QoS settings at the volume level are overridden.
   - 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 **OK**.

Related concepts

Using object naming best practices when managing multiple clusters on page 49

Related tasks

Adding a cluster on page 50
Creating an account on page 109

Viewing volumes details

You can review general information for all active volumes on the cluster in the NetApp Element Management extension point. You can also see details for each active volume, including efficiency, performance, QoS, as well as associated snapshots.

Steps

1. Select **NetApp Element Management** > **Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.

   General information about active volumes is displayed.

3. Select the check box for the individual volume you want to review.

4. Click **Actions**.

5. In the resulting menu, select **View details**.

Volume details

You can view volume details on the **Volumes** page of the **Management** tab from the **NetApp Element Management** extension point.

Note: VMware requires 512e for disk resources. If 512e is not enabled, a VMFS cannot be created.

Volume ID

The system-generated ID for the volume.

Volume Name

The name assigned to the volume.

Account

The name of the account assigned to the volume.

Access Groups

The name of the volume access group 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 is allowed.
- `ReplicationTarget`: Designated as a target volume in a replicated volume pair.

Volume Paired

Indicates whether or not the volume is part of a volume pairing.

Size (GB)

The total size in GB of the volume.

Snapshots

The number of snapshots created for the volume.

QoS Policy

The name of the user-defined QoS policy.

512e

Identifies if 512e is enabled on a volume. The value can be either `Yes` or `No`.

Individual volume details

You can view information for active volumes when you select an individual volume and view its details on the **Volumes** page of the **Management** tab from the **NetApp Element Management** extension point.

These headings are in the **General Details** section:

Name

The name assigned to the volume.

Volume ID

The system-generated ID for the volume.

IQN

The iSCSI Qualified Name of the volume.

Account ID

The unique account ID of the associated account.

Account

The name of the account assigned to the volume.

Access Groups

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

Size

The total size in bytes of the volume.

Volume Paired

Indicates whether or not the volume is part of a volume pairing.

SCSI EUI Device ID

Globally unique SCSI device identifier for the volume in EUI-64 based 16-byte format.

SCSI NAA Device ID

The globally unique SCSI device identifier for the protocol endpoint in NAA IEEE Registered Extended Format.

These headings are in the **Efficiency** section:

Compression

The compression efficiency score for the volume.

Deduplication

The deduplication efficiency score for the volume.

Thin Provisioning

The thin provisioning efficiency score for the volume.

Last Updated

The date and time of the last efficiency score.

These headings are in the **Performance** section:

Account ID

The unique account ID of the associated account.

Actual IOPS

Current actual IOPS to the volume in the last 500 milliseconds.

Async Delay

The length of time since the volume was last synced with the remote cluster.

Average IOP Size

Average size in bytes of recent I/O to the volume in the last 500 milliseconds.

Burst IOPS Size

The total number of IOP credits available to the user. When volumes are not using up to the Max IOPS, credits are accrued.

Client Queue Depth

The number of outstanding read and write operations to the volume.

Last Updated

The date and time of the last performance update.

Latency USec

The average time, in microseconds, to complete operations to the volume in the last 500 milliseconds. A "0" (zero) value means there is no I/O to the volume.

Non-zero Blocks

Total number of 4KiB blocks with data after the last garbage collection operation has completed.

Performance Utilization

The percentage of cluster IOPS being consumed. For example, a 250K IOP cluster running at 100K IOPS would show 40% consumption.

Read Bytes

The total cumulative bytes read from the volume since the creation of the volume.

Read Latency USec

The average time, in microseconds, to complete read operations to the volume in the last 500 milliseconds.

Read Operations

The total read operations to the volume since the creation of the volume.

Thin Provisioning

The thin provisioning efficiency score for the volume.

Throttle

A floating value between 0 and 1 that represents how much the system is throttling clients below their maxIOPS because of re-replication of data, transient errors and snapshots taken.

Total Latency USec

The time, in microseconds, to complete read and write operations to a volume.

Unaligned Reads

For 512e volumes, the number of read operations that were not on a 4k sector boundary. High numbers of unaligned reads may indicate improper partition alignment.

Unaligned Writes

For 512e volumes, the number of write operations that were not on a 4k sector boundary. High numbers of unaligned writes may indicate improper partition alignment.

Used Capacity

Percentage of used capacity.

Volume ID

The system-generated ID for the volume.

Vol Access Groups

The volume access group IDs that are associated with the volume.

Volume Utilization

A percentage value that describes how much the client is using the volume.

Possible values:

- `0`: Client is not using the volume.
- `100`: Client is using their max.
- `>100`: Client is using their burst.

Write Bytes

The total cumulative bytes written to the volume since the creation of the volume.

Write Latency USec

The average time, in microseconds, to complete write operations to a volume in the last 500 milliseconds.

Write Operations

The total cumulative write operations to the volume since the creation of the volume.

Zero Blocks

Total number of 4KiB blocks without data after the last round of garbage collection operation has completed.

These headings are in the **Quality of Service** section:

Policy

The name of the QoS policy assigned to the volume.

I/O Size

The size of the IOPS in KB.

Min IOPS

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

Max IOPS

The maximum number of sustained IOPS that the 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.

Max Bandwidth

The maximum bandwidth permitted by the system to process larger block sizes.

These headings are in the **Snapshots** section:

Snapshot ID

System generated ID for the snapshot.

Snapshot Name

User-defined name for the snapshot.

Create Date

The date and time at which the snapshot was created.

Expiration Date

The day and time the snapshot will be deleted.

Size

User-defined size of the snapshot in GB.

Editing a volume

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

About this task

If you are using persistent volumes with the management node, do not modify the names of the persistent volumes.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for the volume you want to edit.
4. Click **Actions**.
5. In the resulting menu, click **Edit**.
6. Optional: In the **Volume Size** field, enter a different volume size in GB or GiB.

   Note: You can increase, but not decrease, the size of the volume. 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.

7. Optional: Select a different user account.
8. Optional: Select a different account access level of one of the following:
   - **Read/Write**
   - **Read Only**
   - **Locked**
   - **Replication Target**
9. In the **Quality of Service** area, do one of the following:
   - Under **Policy**, select an existing QoS policy, if available.
   - Under **Custom Settings**, set customized minimum, maximum, and burst values for IOPS or use the default QoS values.

Best Practice: When you change IOPS values, use increments in tens or hundreds. Input values require valid whole numbers.

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.

Attention:

- Datastore QoSSIOC settings will override any QoS settings at the volume level.
- 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.

10. Click **OK**.

Cloning a volume

You can create a clone of a volume 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.

Before you begin

- At least one cluster must be added and running.
- At least one volume must be created.
- A user account has been created.
- Available unprovisioned space must be equal to or more than the volume size.

About this task

The cluster supports up to two running clone requests per volume at a time and up to 8 active volume clone operations at a time. Requests beyond these limits are queued for later processing.

Note: Cloned volumes do not inherit volume access group membership from the source volume.

Steps

1. Select **NetApp Element Management > Management.**

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the Active view, select the check box for the volume you want to clone.
4. Click **Actions**.
5. In the resulting menu, click **Clone**.
6. In the **Clone Volume** dialog box, enter a volume name for the newly cloned volume.

   Tip: Use descriptive naming best practices. This is especially important if multiple clusters or vCenter Servers are used in your environment.

7. Select a size in GB or GIB for the cloned 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

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 may need to extend partitions or create new partitions in the free space to make use of it.

8. Select an account to associate with the newly cloned volume.
9. Select the one of the following access types for the newly cloned volume:
   - **Read/Write**
   - **Read Only**
   - **Locked**
10. Adjust 512e settings, if required.

    Note:

    - By default, 512 byte emulation is enabled for all new volumes.
    - VMware requires 512e for disk resources. If 512e is not enabled, a VMFS cannot be created and volume details are grayed out.

11. Click **OK**.

    Note: The time to complete a cloning operation is affected by volume size and current cluster load. Refresh the page if the cloned volume does not appear in the volume list.

Related concepts

Using object naming best practices when managing multiple clusters on page 49

Related tasks

Adding a cluster on page 50
Creating an account on page 109
Creating a volume on page 91

Volume backup and restore operations

You can configure the system to back up and restore the contents of a volume to and from an object store container that is external to NetApp Element software-based storage.

You can also back up and restore data to and from remote NetApp Element software-based systems. You can run a maximum of two backup or restore processes at a time on a volume.

Volume backup operations

You can back up NetApp Element volumes to Element storage, as well as secondary object stores that are compatible with Amazon S3 or OpenStack Swift.

Related tasks

Backing up a volume to an Amazon S3 object store on page 100

Backing up a volume to an OpenStack Swift object store on page 100
Backing up a volume to a cluster running Element software on page 101

Backing up a volume to an Amazon S3 object store

You can back up NetApp Element volumes to external object stores that are compatible with Amazon S3.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for the volume you want to back up.
4. Click **Actions**.
5. In the resulting menu, click **Back Up to**.
6. In the dialog box under **Back up volume to**, select **Amazon S3**.
7. Select an option under with the following data format:
   - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
   - **Uncompressed**: An uncompressed format compatible with other systems.
8. Enter a host name to use to access the object store in the **Host name** field.
9. Enter an access key ID for the account in the **Access key ID** field.
10. Enter the secret access key for the account in the **Secret access key** field.
11. Enter the S3 bucket in which to store the backup in the **Amazon S3 bucket** field.
12. Optional: Enter a prefix for the backup volume name in the **Prefix** field.
13. Optional: Enter a nametag to append to the prefix in the **Nametag** field.
14. Click **OK**.

Backing up a volume to an OpenStack Swift object store

You can back up NetApp Element volumes to external object stores that are compatible with OpenStack Swift.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for the volume you want to back up.
4. Click **Actions**.

5. In the resulting menu, click **Back Up to**.
6. In the dialog box under **Back up volume to**, select **OpenStack Swift**.
7. Select an option under **with the following data format**:
   - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
   - **Uncompressed**: An uncompressed format compatible with other systems.
8. Enter a URL to use to access the object store in the **URL** field.
9. Enter a user name for the account in the **User name** field.
10. Enter the authentication key for the account in the **Authentication key** field.
11. Enter the container in which to store the backup in the **Container** field.
12. Optional: Enter a prefix for the backup volume name in the **Prefix** field.
13. Optional: Enter a nametag to append to the prefix in the **Nametag** field.
14. Click **OK**.

Backing up a volume to a cluster running Element software

You can back up volumes residing on a cluster running NetApp Element software to a remote Element cluster. 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 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. From the vCenter that contains the destination cluster, select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for the destination volume.
4. Click **Actions**.
5. In the resulting menu, click **Restore from**.
6. In the dialog box under **Restore from**, select **NetApp Element**.
7. Select an option under **with the following data format**:
   - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
   - **Uncompressed**: An uncompressed format compatible with other systems.
8. Click **Generate Key** to generate a bulk volume write key for the destination volume.
9. Copy the bulk volume write key to your clipboard to apply to later steps on the source cluster.

10. From the vCenter that contains the source cluster, select **NetApp Element Management** > **Management**.
11. Click the **Volumes** sub-tab.
12. From the **Active** view, select the check box for the destination volume.
13. Click **Actions**.
14. In the resulting menu, click **Back Up to**.
15. In the dialog box under **Back up volume to**, select **NetApp Element**.
16. Select the same option as the destination cluster under **with the following data format**:
17. Enter the management virtual IP address of the destination volume's cluster in the **Remote cluster MVIP** field.
18. Enter the cluster administrator user name for the destination cluster in the **Remote cluster user name** field.
19. Enter the cluster administrator password for the destination cluster in the **Remote cluster user password** field.
20. In the **Bulk volume write key** field, paste the key you generated on the destination cluster.
21. Click **OK**.

Volume restore operations

When you restore a volume from a backup on an object store such as OpenStack Swift or Amazon S3, you need manifest information from the original backup process. If you are restoring a NetApp Element volume that was backed up on a NetApp Element-based storage system, the manifest information is not required. You can find the required manifest information for restoring from Swift and S3 in the **Event Log** on the **Reporting** tab.

Related tasks

Restoring a volume from backup on an Amazon S3 object store on page 102
Restoring a volume from backup on an OpenStack Swift object store on page 103
Restoring a volume from backup on a cluster running Element software on page 104

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 using the plug-in.

Steps

1. Select **NetApp Element Management** > **Reporting**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Event Log** sub-tab.
3. Select the backup event that created the backup you need to restore.
4. Click the **Details** button for the event.
5. In the resulting menu, click **View Details**.
6. Copy the manifest information to your clipboard.

7. Click **Management > Volumes**.
8. From the **Active** view, select the check box for the volume you want to restore.
9. Click **Actions**.
10. In the resulting menu, click **Restore from**.
11. In the **Restore Volume** dialog box under **Restore from**, select **Amazon S3**.
12. Select an option under **with the following data format**:
    - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
    - **Uncompressed**: An uncompressed format compatible with other systems.
13. Enter a host name to use to access the object store in the **Host name** field.
14. Enter an access key ID for the account in the **Access key ID** field.
15. Enter the secret access key for the account in the **Secret access key** field.
16. Enter the S3 bucket in which to store the backup in the **Amazon S3 bucket** field.
17. Paste the manifest information into the **Manifest** field.
18. Click **OK**.

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 using the plug-in.

Steps

1. Select **NetApp Element Management > Reporting**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Event Log** sub-tab.
3. Select the backup event that created the backup you need to restore.
4. Click the **Details** button for the event.
5. In the resulting menu, click **View Details**.
6. Copy the manifest information to your clipboard.
7. Click **Management > Volumes**.
8. From the **Active** view, select the check box for the volume you want to restore.
9. Click **Actions**.
10. In the resulting menu, click **Restore from**.
11. In the **Restore Volume** dialog box under **Restore from**, select **OpenStack Swift**.
12. Select an option under **with the following data format**:
    - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
    - **Uncompressed**: An uncompressed format compatible with other systems.

13. Enter a URL to use to access the object store in the **URL** field.
14. Enter a user name for the account in the **User name** field.
15. Enter the authentication key for the account in the **Authentication key** field.
16. Enter the name of the container in which the backup is stored in the **Container** field.
17. Paste the manifest information into the **Manifest** field.
18. Click **OK**.

Restoring a volume from backup on a cluster running Element software

You can restore a volume from a backup on a cluster running NetApp Element software. 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 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. From the vCenter that contains the destination cluster, select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for the volume you want to restore.
4. Click **Actions**.
5. In the resulting menu, click **Restore from**.
6. In the **Restore Volume** dialog box under **Restore from**, select **NetApp Element**.
7. Select an option under **with the following data format**:
   - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
   - **Uncompressed**: An uncompressed format compatible with other systems.
8. Click **Generate Key** to generate a bulk volume write key for the destination volume.
9. Copy the bulk volume write key to your clipboard to apply to later steps on the source cluster.
10. From the vCenter that contains the source cluster, select **NetApp Element Management > Management**.
11. Click the **Volumes** sub-tab.
12. From the **Active** view, select the check box for the volume you want to use for the restore.
13. Click **Actions**.
14. In the resulting menu, click **Backup to**.
15. In the dialog box under **Back up volume to**, select **NetApp Element**.
16. Select the option that matches the backup under **with the following data format**.

17. Enter the management virtual IP address of the destination volume's cluster in the **Remote cluster MVIP** field.
18. Enter the cluster administrator user name for the destination cluster in the **Remote cluster user name** field.
19. Enter the cluster administrator password for the destination cluster in the **Remote cluster user password** field.
20. In the **Bulk volume write key** field, paste the key you generated on the destination cluster.
21. Click **OK**.

Deleting volumes

You can delete one or more volumes from a NetApp Element cluster using the NetApp Element Management extension point.

About this task

The system does not immediately purge a deleted volume. A deleted volume can be restored for approximately eight hours. You can restore a volume before the system purges it or manually purge the volume from the **Deleted** view in **Management > Volumes**. When you restore a volume, it comes back online and iSCSI connections are restored.

Attention: If a volume used to create a snapshot is deleted, its associated snapshots are listed in the **Inactive** view on the **Protection > Snapshots** page. When the deleted source volumes are purged, the snapshots in **Inactive** view are also removed from the system.

Note: Do not delete persistent volumes that have been created by the management node, if applicable.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. Delete one or more volumes:
   a. From the **Active** view, select the check box for each volume you want to delete.
   b. Click **Actions**.
   c. In the resulting menu, click **Delete**.

      Note: The plug-in does not allow a volume with a datastore to be deleted.

   d. Confirm the action.

      The system moves the volume from the **Active** view to the **Deleted** view in the **Volumes** page.

Related tasks

Restoring deleted volumes on page 106

Purging volumes

You can manually purge volumes after you have deleted them.

About this task

The system automatically purges deleted volumes eight hours after deletion. However, if you want to purge a volume before the scheduled purge time, you can perform a manual purge using the following steps.

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

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. Select the view filter and select **Deleted** from the drop-down list.
4. Select one or more volumes you want to purge.
5. Click **Purge**.
6. Confirm the action.

Restoring deleted volumes

You can restore a volume in the NetApp Element system if it has been deleted but not yet purged.

About this task

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.

Note: If a volume is deleted and then restored, ESXi will not detect the restored volume (and datastore if it exists). Remove the static target from the ESXi iSCSI adapter and rescan the adapter.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. Select the view filter and select **Deleted** from the drop-down list.
4. Select one or more volumes you want to restore.
5. Click **Restore**.
6. Select the view filter and select **Active** from the drop-down list.
7. Verify that the volume or volumes and all connections are restored.

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 from the **Active** volumes view.

Before you begin

- At least one cluster must be added and running.
- At least one access group exists.
- At least one active volume exists.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. Select the check box for each volume you want to add to an access group.
4. Click **Actions**.
5. In the resulting menu, select **Add to Access Group**.
6. Confirm the details in the dialog box that opens, and select a volume access group from the drop-down list.
7. Click **OK**.

Related tasks

Creating a volume on page 91
Creating access groups on page 111

Removing volumes from an access group

You can remove volumes from an access group.

About this task

When you remove a volume from an access group, the group no longer has access to that volume.

Attention: Removing a volume from an access group can disrupt host access to the volume.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. Select the check box for each volume you want to remove from an access group.
4. Click **Actions**.
5. In the resulting menu, select **Remove from Access Group**.

6. Confirm the details in the dialog box that opens, and select the volume access group that you no longer want to have access to each selected volume.
7. Click **OK**.

Applying a QoS policy to multiple volumes

You can apply an existing QoS policy to multiple volumes. Use this process when you want to bulk apply a policy to one or more volumes.

Before you begin

The QoS policy you want to bulk apply exists.

Steps

1. Select **NetApp Element Management** > **Management**.
2. Click the **Volumes** sub-tab.
3. Select the check box for each volume to which you want to apply a QoS policy.
4. Click **Actions**.
5. In the resulting menu, select **Apply QoS Policy**.
6. In the dialog box, select the QoS policy from the drop-down list to apply to the selected volumes.
7. Click **OK**.

Related tasks

Creating a QoS policy on page 116

Changing the QoS policy association of a volume

You can remove a QoS policy association from a volume or select a different QoS policy or custom QoS using the NetApp Element Management extension point.

Before you begin

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

Steps

1. Select **NetApp Element Management** > **Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. Select the check box for a volume that contains a QoS policy you want to modify.
4. Click **Actions**.
5. In the resulting menu, select **Edit**.
6. In the dialog box under **Quality of Service**, select a new QoS policy or custom settings to apply to the volume.
7. If you chose **Custom Settings**, modify the **Min IOPS**, **Max IOPS**, and **Burst IOPS** values.

Note: You can also click **Reset Default QoS** to restore default IOPS values.

8. Click **OK**.

User account management

User accounts are used to control access to the storage resources on a NetApp Element software-based network. At least one user account is required before a volume can be created. When you create a volume, it is assigned to an account. If you have created a virtual volume, the account is the storage container. The 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.

User accounts can be managed from NetApp Element Management extension point.

Creating an account

You can create an account to allow access to volumes. After you create an account, you can assign up to 2000 volumes to the account. Each account name in the system must be unique.

Before you begin

At least one cluster must be added and running.

Steps

1. Select **NetApp Element Management** > **Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Accounts** sub-tab.
3. Click **Create Account**.
4. Enter a user name.

   Tip: Use descriptive naming best practices. This is especially important if multiple clusters or vCenter Servers are used in your environment.

5. In the **CHAP Settings** section:
   a. Enter the initiator secret for CHAP node session authentication.
   b. Enter the target secret for CHAP node session authentication.

      Note: Initiator and target secrets must differ. If these fields are left blank, the system generates the authentication credentials.

6. Click **OK**.

Account details

You can view a list of all accounts on the **Accounts** page of the **Management** tab from the **NetApp Element Management** extension point.

Account ID

System-generated ID for the account.

User name

The name given to the account when it was created.

Number of Volumes

The number of active volumes assigned to the account.

Status

The status of the account.

Editing an account

You can edit an account to change the status or the CHAP secrets. Changing CHAP settings can cause lost connectivity between a host and its associated volumes.

About this task

If you are using persistent volumes with the management node, do not modify the account name of the account associated with these volumes.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Accounts** sub-tab.
3. Select the check box for the account you want to edit.
4. Click **Actions**.
5. In the resulting menu, click **Edit**.
6. Optional: Under **Edit Account**, edit the access status of the account.

   Attention: Changing the access 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.

7. Optional: Edit the initiator secret or target secret credentials used for node session authentication.

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

8. Click **OK**.

Deleting an account

You can delete accounts using the NetApp Element Management extension point after volumes associations with the account have been removed.

Before you begin

Delete and purge any volumes associated with the account or reassign the volumes to another account.

About this task

If you are using persistent volumes with the management node, do not delete the account associated with these volumes.

Steps

1. Select **NetApp Element Management > Management**.

Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Accounts** sub-tab.
3. Select the check box for the account you want to delete.
4. Click **Actions**.
5. In the resulting menu, select **Delete**.
6. Confirm the action.

Related tasks

Deleting volumes on page 105
Purging volumes on page 106

Volume access groups

A volume access group is a collection of volumes that users can access using either iSCSI initiators or Fibre Channel initiators.

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.

Volume access groups can be managed from the NetApp Element Management extension point.

Creating access groups

You can create volume access groups with one or multiple initiators. Mapping Fibre Channel (WWPN) or iSCSI (IQN) client initiators to the volumes in a volume access group enables secure data I/O between a network and a volume.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Access Groups** sub-tab.
3. Click **Create Access Group**.
4. Enter a name for the volume access group.

   Tip: Use descriptive naming best practices. This is especially important if multiple clusters or vCenter Servers are used in your environment.

5. Select an unassigned IQN or WWPN from the **Select an Initiator** drop-down list and click **Add Initiator**.

   Note: Initiators may be added or deleted after the volume access group has been created.

6. Click **OK**.

Volume access group details

You can view volume access group information on the **Access Groups** page of the **Management** tab from the **NetApp Element Management** extension point.

ID

System-generated ID for the volume access group.

Name

The name given to the volume access group when it was created.

Active Volumes

The number of active volumes in the volume access group.

Initiators

The number of initiators connected to the volume access group.

Editing access groups

You can edit volume access group names or add or remove initiators from the NetApp Element Management extension point.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Access Groups** sub-tab.
3. Select the check box for the volume access group you want to edit.
4. Click **Actions**.
5. In the resulting menu, select **Edit**.
6. Optional: Modify the access group name.
7. Optional: Add or remove initiators.

   Note: If you are removing an initiator, click the trash icon to remove it. When you remove the initiator, it can no longer access the volumes in that volume access group. Normal account access to the volume is not disrupted.

8. Click **OK**.

Deleting access groups

You can delete volume access groups using the NetApp Element Management extension point. You do not need to delete initiator IDs or disassociate volumes from the volume access group prior to deleting the group. After you delete the access group, group access to the volumes is discontinued.

Steps

1. Select **NetApp Element Management > Management**

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Access Groups** sub-tab.
3. Select the check box for the volume access group you want to delete.
4. Click **Actions**.
5. In the resulting menu, select **Delete**.
6. Confirm the action.

Initiators

Initiators enable external clients access to volumes in a cluster, serving as the entry point for communication between clients and volumes. You can create and delete initiators, and give them friendly aliases to simplify administration and volume access. When you add an initiator to a volume access group, that initiator enables access to all volumes in the group.

You can view initiators on the **Management > Initiators** page from the NetApp Element Management extension point.

Creating an initiator

You can create iSCSI or Fibre Channel initiators and optionally assign them aliases.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Initiators** sub-tab.
3. Click **Create Initiator**.
4. To create a single initiator:

   a. Select **Create a Single Initiator**.

   b. Enter the IQN or WWPN for the initiator in the **IQN/WWPN** field.

   - The accepted format of an initiator IQN is `iqn.yyyy-mm` where `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
```

- The accepted format of a Fibre Channel initiator WWPN is `:Aa:bB:CC:dd:11:22:33:44` or `AabBCCdd11223344`.
  A sample of the format is as follows:

```
5f:47:ac:c0:5c:74:d4:02
```

c. Enter a friendly name for the initiator in the **Alias** field.

5. To create multiple initiators:

   a. Select **Create Multiple Initiators**.

   b. Do one of the following:

   - Click **Scan Hosts** to scan vSphere hosts for initiator values not defined in the NetApp Element cluster.
   - Enter a list of IQNs or WWPNs in the text box, and click **Add Initiators**.

   c. Optional: Under the **Alias** heading, click the field for each entry to add an alias.

   d. Optional: Remove an initiator from the list, as required.

6. Click **OK**.

Initiator details

You can view information about initiators on the **Initiators** page of the **Management** tab from the **NetApp Element Management** extension point.

ID
The system-generated ID for the initiator.

Name
The name given to the initiator when it was created.

Alias
The friendly name given to the initiator, if any.

Access Group
The volume access group to which the initiator is assigned.

Editing an initiator

You can change the alias of an existing initiator or add an alias if one does not already exist.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Initiators** sub-tab.
3. Select the check box for the initiator you want to edit.
4. Click **Actions**.

5. In the resulting menu, click **Edit**.
6. Enter a new alias for the initiator in the **Alias** field.
7. Click **OK**.

Deleting initiators

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. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Initiators** sub-tab.
3. Select the check boxes for the initiators you want to delete.
4. Click **Actions**.
5. In the resulting menu, select **Delete**.
6. Confirm the action.

Adding initiators to a volume access group

You can add initiators to an access group to allow access to volumes in the volume access group without requiring CHAP authentication. When you add an initiator to a volume access group, the initiator has access to all volumes in that volume access group.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Initiators** sub-tab.
3. Select the check boxes for the initiators you want to add to an access group.
4. Click **Actions**.
5. In the resulting menu, click **Add to Access Group**.
6. In the **Add to Access Group** dialog, choose an access group from the drop-down list.
7. Click **OK**.

QoS policies

A QoS (Quality of Service) 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.

You can view QoS policies on the **Management > QoS Policies** page from the NetApp Element Management extension point.

Note: If you are using QoS policies, do not enable QoSSIOC. QoSSIOC will override any volume QoS settings.

Note: The selected cluster must be Element 10.0 or later to use QoS policies; otherwise, QoS policy functions are not available.

Creating a QoS policy

You can create QoS policies and apply them when creating volumes from the NetApp Element Management extension point.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **QoS Policies** sub-tab.
3. Click **Create QoS Policy**.
4. Enter the **Policy Name**.

   Tip: Use descriptive naming best practices. This is especially important if multiple clusters or vCenter Servers are used in your environment.

5. Enter the **Min IOPS**, **Max IOPS**, and **Burst IOPS** values.
6. Click **OK**.

QoS policy details

You can view information about QoS policies on the **QoS Policies** page of the **Management** tab from the **NetApp Element Management** extension point.

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 from the NetApp Element Management extension point. Changing a QoS policy affects all volumes associated with the policy.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **QoS Policies** sub-tab.
3. Select the check box for the QoS policy you want to edit.
4. Click **Actions**.
5. In the resulting menu, select **Edit**.
6. In the **Edit QoS Policy** dialog box, modify the following properties as needed:
   - **Policy Name**
   - **Min IOPS**
   - **Max IOPS**
   - **Burst IOPS**

   Note: You can also click **Reset Default QoS** to restore default IOPS values.

7. Click **OK**.

Deleting a QoS policy

You can delete a QoS policy if it is no longer needed from the NetApp Element Management extension point. 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 **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **QoS Policies** sub-tab.
3. Select the check box for the QoS policy you want to delete.
4. Click **Actions**.
5. In the resulting menu, select **Delete**.
6. Confirm the action.

Data protection

From the **Protection** tab, you can perform tasks that ensure copies of your data are created and stored where you need them. These tasks include creating and managing volume and group snapshots, snapshot schedules, and creating volume and cluster pair relationships for replication between remote clusters.

Volume snapshots

A volume snapshot is a point-in-time copy of a volume. Creating a volume snapshot takes only a small amount of system resources and space; this makes snapshot creation faster than cloning. You can use snapshots to roll a volume back to the state it was in at the time the snapshot was created. However, because snapshots are simply replicas of volume metadata, you cannot mount or write to them.

You can replicate snapshots to a remote NetApp Element cluster and use them as a backup copy for 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.

You can create volume snapshots from the **NetApp Element Management > Management > Volumes** page. You can manage these volume snapshots from the **NetApp Element Management > Protection > Snapshots** page.

Creating a volume snapshot

You can create a snapshot of the active volume to preserve the volume image at any point in time. You can create the snapshot immediately or create a schedule to automate future snapshots of the volume. You can create up to 32 snapshots for a single volume.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for the volume you want to use for the snapshot.
4. Click **Actions**.
5. In the resulting menu, select **Create Snapshot**.
6. Optional: In the **Create Snapshot** dialog box, enter a name for the snapshot.

   Tip: Use descriptive naming best practices. This is especially important if multiple clusters or vCenter Servers are used in your environment.

   Note: If you do not enter a name, the system creates a snapshot default name using the date and time the snapshot was created.

7. Optional: Select the **Include snapshot in replication when paired** check box to ensure that the snapshot is replicated when the parent volume is paired.
8. Select one of the following as the retention period for the snapshot:

- **Keep forever**: Retains the snapshot on the system indefinitely.
- **Set retention period**: Determine a length of time (days, hours, or minutes) for the system to retain the snapshot.

  Note: When you set a retention period, you select a period that begins at the current time (retention is not calculated from the snapshot creation time).

9. To take a single, immediate snapshot, select **Take snapshot now**.
10. To schedule the snapshot to run at a future time, perform the following steps:
    a. Select **Create snapshot schedule**.
    b. Enter a schedule name.
    c. Select a schedule type from the list and configure the schedule details.
    d. Optional: Select the **Recurrent Schedule** check box to repeat the scheduled snapshot periodically.
11. Click **OK**.

Volume snapshot details

You can view information about volume snapshots on a cluster on the **Snapshots** page of the **Protection** tab from the **NetApp Element Management** extension point.

You can also filter snapshots by the following views:

- **Individual**: Volume snapshots that are not members of a group snapshot.
- **Members**: Volume snapshots that are members of a group snapshot.
- **Inactive**: Volume snapshots that were created from volumes that have been deleted but not yet purged.

ID

System generated ID for the snapshot.

Snapshot UUID

The unique ID of the snapshot.

Name

User-defined name or system default name for the snapshot.

Size

User-defined size of the snapshot.

Volume ID

ID of the volume from which the snapshot was created.

Volume Name

User defined name of the volume.

Account

Account associated with the volume.

Volume Size

Size of the volume from which the snapshot was created.

Create Date

The date and time at which the snapshot was created.

Expiration Date

The day and time the snapshot will be deleted as defined by the retention period.

Group Snapshot ID

The group ID the snapshot belongs to if grouped together with other volume snapshots.

Remote Replication

Identifies whether or not the snapshot is enabled for replication to a remote cluster running NetApp Element software.

Possible Values:

- `True`: The snapshot is enabled for remote replication.
- `False`: The snapshot is not enabled for remote replication.

Remote Status

Displays the status of the snapshot on the remote cluster running NetApp Element software.

Possible Values:

- `Present`: The snapshot exists on a remote cluster.
- `Not Present`: The snapshot does not exist on a remote cluster.
- `Syncing`: The target cluster is currently replicating the snapshot.
- `Deleted`: The target replicated the snapshot and then deleted it.

Editing snapshots

You can change replication settings or the retention period for a snapshot. 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. Select **NetApp Element Management** > **Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Snapshots** sub-tab, select one of two views:
   - **Individual**: Volume snapshots that are not members of a group snapshot.
   - **Members**: Volume snapshots that are members of a group snapshot.
3. Select the check box for the volume snapshot you want to edit.
4. Click **Actions**.
5. In the resulting menu, select **Edit**.
6. 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.
7. Optional: Select one of the following as the retention period for the snapshot:
   - **Keep forever**: Retains the snapshot on the system indefinitely.

- **Set retention period**: Determine a length of time (days, hours, or minutes) for the system to retain the snapshot.

  Note: When you set a retention period, you select a period that begins at the current time (retention is not calculated from the snapshot creation time).

8. Click **OK**.

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 also stores information about other snapshots of the volume in the new created volume.

Steps

1. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Snapshots** sub-tab, select one of two views:

   - **Individual**: Volume snapshots that are not members of a group snapshot.
   - **Members**: Volume snapshots that are members of a group snapshot.

3. Select the check box for the volume snapshot you want to clone as a volume.
4. Click **Actions**.
5. In the resulting menu, click **Clone Volume from Snapshot**.
6. In the **Clone Volume from Snapshot** dialog box, enter a volume name.
7. Enter the total size and select either GB or GiB for the new volume.
8. Select an access type for the volume:

   - **Read Only**: Only read operations are allowed.
   - **Read / Write**: Reads and writes are allowed.
   - **Locked**: No reads or writes are allowed.
   - **Replication Target**: Designated as a target volume in a replicated volume pair.

9. Select a user account from the list to associate with the new volume.
10. Click **OK**.
11. Select **NetApp Element Management > Management**.
12. Click the **Volumes** sub-tab.
13. From the **Active** view, confirm that the new volume is listed.

    Note: Refresh the page if needed until the new volume appears in the list.

Rolling back a volume to a snapshot

You can roll back a volume to a snapshot at any time. This undoes any changes made to the volume since the snapshot was created.

Steps

1. Select **NetApp Element Management** > **Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Snapshots** sub-tab, select one of two views:

   - **Individual**: Volume snapshots that are not members of a group snapshot.
   - **Members**: Volume snapshots that are members of a group snapshot.

3. Select the check box for the volume snapshot you want to use for the volume rollback.
4. Click **Actions**.
5. In the resulting menu, select **Rollback Volume to Snapshot**.
6. 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.

7. Click **OK**.

Volume snapshot backup operations

You can use the integrated backup feature to back up a volume snapshot. You can back up snapshots from a cluster running NetApp Element software to an external object store or to another Element-based 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.

Related tasks

Backing up a volume snapshot to an Amazon S3 object store on page 122
Backing up a volume snapshot to an OpenStack Swift object store on page 123
Backing up a volume snapshot to a cluster running Element software on page 124

Backing up a volume snapshot to an Amazon S3 object store

You can back up NetApp Element snapshots to external object stores that are compatible with Amazon S3.

Steps

1. Select **NetApp Element Management** > **Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Snapshots** sub-tab, select the check box for the volume snapshot you want to back up.
3. Click **Actions**.

4. In the resulting menu, click **Backup to**.
5. In the dialog under **Back up volume to**, select **Amazon S3**.
6. Select an option under **with the following data format**:
   - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
   - **Uncompressed**: An uncompressed format compatible with other systems.
7. Enter a host name to use to access the object store in the **Host name** field.
8. Enter an access key ID for the account in the **Access key ID** field.
9. Enter the secret access key for the account in the **Secret access key** field.
10. Enter the S3 bucket in which to store the backup in the **Amazon S3 Bucket** field.
11. Optional: Enter a prefix for the backup name in the **Prefix** field.
12. Optional: Enter a nametag to append to the prefix in the **Nametag** field.
13. Click **OK**.

Backing up a volume snapshot to an OpenStack Swift object store

You can back up NetApp Element snapshots to secondary object stores that are compatible with OpenStack Swift.

Steps

1. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Snapshots** sub-tab, select the check box for the volume snapshot you want to back up.
3. Click **Actions**.
4. In the resulting menu, click **Backup to**.
5. In the dialog box under **Back up volume to**, select **OpenStack Swift**.
6. Select an option under **with the following data format**:
   - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
   - **Uncompressed**: An uncompressed format compatible with other systems.
7. Enter a URL to use to access the object store in the **URL** field.
8. Enter a user name for the account in the **User name** field.
9. Enter the authentication key for the account in the **Authentication key** field.
10. Enter the container in which to store the backup in the **Container** field.
11. Optional: Enter a prefix for the backup volume name in the **Prefix** field.
12. Optional: Enter a name tag to append to the prefix in the **Nametag** field.
13. Click **OK**.

Backing up a volume snapshot to a cluster running Element software

You can back up a volume snapshot that resides on a cluster running NetApp Element software to a remote Element cluster.

Before you begin

You must create a volume on the destination cluster of equal or greater size to the snapshot you are using for the backup.

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 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. From the vCenter that contains the destination cluster, select **NetApp Element Management** > **Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Volumes** sub-tab, select the check box for the destination volume.
3. Click **Actions**.
4. In the resulting menu, click **Restore from**.
5. In the dialog box under **Restore from**, select **NetApp Element**.
6. Select an option under **with the following data format**:
   - **Native**: A compressed format readable only by NetApp Element software-based storage systems.
   - **Uncompressed**: An uncompressed format compatible with other systems.
7. Click **Generate Key** to generate a bulk volume write key for the destination volume.
8. Copy the bulk volume write key to your clipboard to apply to later steps on the source cluster.
9. From the vCenter that contains the source cluster, select **NetApp Element Management** > **Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

10. Select the check box for the snapshot you are using for the backup.
11. Click **Actions**.
12. In the resulting menu, click **Backup to**.
13. In the dialog box under **Back up volume to**, select **NetApp Element**.
14. Select the same option as the destination cluster under **with the following data format**:

15. Enter the management virtual IP address of the destination volume's cluster in the **Remote cluster MVIP** field.
16. Enter the remote cluster user name in the **Remote cluster user name** field.
17. Enter the remote cluster password in the **Remote cluster user password** field.
18. In the **Bulk volume write key** field, paste the key you generated on the destination cluster earlier.
19. Click **OK**.

Deleting a snapshot

You can delete a volume snapshot from a cluster running NetApp Element software using the NetApp Element Management extension point. 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. After the target has detected that you have deleted the source snapshot, the target stops replication of the snapshot.

Steps

1. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Snapshots** sub-tab, select one of two views:

   - **Individual**: A list of volume snapshots that are not part of a group snapshot.
   - **Inactive**: A list of volume snapshots that were created from volumes that have been deleted but not yet purged.

   Note: You cannot delete individual members of a group snapshot from the **Snapshots** sub-tab. You must instead delete the group snapshot. During group snapshot deletion, you are given the option to convert members of a group snapshot to individual (group snapshot unaffiliated) snapshots.

3. Select the check box for the volume snapshot you want to delete.
4. Click **Actions**.
5. In the resulting menu, select **Delete**.
6. Confirm the action.

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 as a backup or rollback to restore the state of the group of volumes to a desired point in time.

Creating a group snapshot

You can create a snapshot of a group of volumes immediately or create a schedule to automate future snapshots of the group of volumes. A single group snapshot can consistently snapshot up to 32 volumes at one time.

Steps

1. Select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for each volume you want to use for the snapshot.
4. Click **Actions**.
5. In the resulting menu, select **Create Group Snapshot**.
6. Optional: In the **Create Group Snapshot** dialog box, enter a name for the group snapshot.

   Tip: Use descriptive naming best practices. This is especially important if multiple clusters or vCenter Servers are used in your environment.

   Note: If you do not enter a name, the system creates a group snapshot default name using the date and time the snapshot was created.

7. Optional: Select the **Include snapshot in replication when paired** check box to ensure that the snapshot is replicated when a parent volume is paired.
8. Select one of the following as the retention period for the snapshot:

   - **Keep forever**: Retains the snapshot on the system indefinitely.
   - **Set retention period**: Determine a length of time (days, hours, or minutes) for the system to retain the snapshot.

   Note: When you set a retention period, you select a period that begins at the current time (retention is not calculated from the snapshot creation time).

9. To take a single, immediate snapshot, select **Take group snapshot now**.
10. To schedule the snapshot to run at a future time:

    a. Select **Create snapshot schedule**.

    b. Enter a schedule name.

    c. Select a schedule type from the list and configure the schedule details.

    d. Optional: Select the **Recurrent Schedule** check box to repeat the scheduled snapshot periodically.

11. Click **OK**.

Group snapshot details

You can view information about group volume snapshots on the **Group Snapshots** page of the **Protection** tab from the **NetApp Element Management** extension point.

Snapshot Group ID

System-generated ID for the group snapshot.

Unique ID

The unique ID of the group snapshot.

Snapshot Group Name

User-defined name or system default name for the group snapshot.

Create Date

The date and 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.

Number of Volumes

The number of volumes in the group snapshot.

Editing group snapshots

You can change replication settings or the retention period for a group snapshot. 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. Select **NetApp Element Management** > **Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Group Snapshots** sub-tab.
3. Select the check box for the group snapshot you want to edit.
4. Click **Actions**.
5. In the resulting menu, select **Edit**.
6. Optional: Select the **Include snapshot in replication when paired** check box to ensure that the snapshot is replicated when the parent volume is paired.
7. Optional: Select one of the following as the retention period for the snapshot:

   - `Keep forever`: Retains the snapshot on the system indefinitely.

- `Set retention period`: Determine a length of time (days, hours, or minutes) for the system to retain the snapshot.

  Note: When you set a retention period, you select a period that begins at the current time (retention is not calculated from the snapshot creation time).

8. Click **OK**.

Cloning volumes from a group snapshot

You can clone a group of volumes from a point-in-time group snapshot. After you create the volumes, you can use them like any other volume in the system.

Steps

1. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Group Snapshots** sub-tab.
3. Select the check box for the group snapshot you want to use for the volume clones.
4. Click **Actions**.
5. In the resulting menu, select **Clone Volumes from Group Snapshot**.
6. Optional: In the **Clone Volumes From Group Snapshot** dialog box, enter a new volume name prefix.

   Note: The prefix distinguishes the new clone volumes from existing volumes. The prefix is applied to all volumes created from the group snapshot.

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

   Note: Volume size and current cluster load affect the time needed to complete a cloning operation.

Rolling back volumes to a group snapshot

You can roll back a group of active volumes to a group snapshot. This restores all the associated volumes in a group snapshot to their state at the time the group snapshot was created. This procedure 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. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Group Snapshots** sub-tab.
3. Select the check box for the group snapshot you wish to use for the volume clones.
4. Click **Actions**.
5. In the resulting menu, select **Rollback Volumes to Group Snapshot**.
6. 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 group snapshot.

7. Click **OK**.

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. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Group Snapshots** sub-tab.
3. Select the check box for the group snapshot you want to delete.
4. Click **Actions**.
5. In the resulting menu, click **Delete**.
6. Do one of the following in the confirmation dialog box:

   - Select **Delete group snapshot and all group snapshot members** to delete the group snapshot and all member snapshots.
   - Select **Retain group snapshot members as individual snapshots** to delete the group snapshot but keep all member snapshots.

7. Confirm the action.

Snapshot schedules

You can schedule a snapshot of a volume to automatically occur at specified date and time intervals. You can schedule either single volume snapshots or group snapshots to run automatically.

When you create snapshot schedules, you can store the resulting snapshots on a remote NetApp Element storage system if the volume is being replicated.

Note: Schedules are created using UTC+0 time. You may need to adjust the actual time a snapshot will run based on your time zone.

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 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. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Schedules** sub-tab.
3. Click **Create Schedule**.
4. In the **Create Schedule** dialog box 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.
5. Enter a new schedule name.
6. Select a schedule type from the list and configure the schedule details.
7. Optional: Select **Recurring Schedule** to repeat the snapshot schedule indefinitely.
8. Optional: Enter a name for the new snapshot in the **New Snapshot Name** field.

   Note: If you do not enter a name, the system creates a snapshot default name using the date and time the snapshot was created.

9. Optional: Select the **Include snapshot in replication when paired** check box to ensure that the snapshot is replicated when the parent volume is paired.
10. Select one of the following as the retention period for the snapshot:

    - **Keep forever**: Retains the snapshot on the system indefinitely.
    - **Set retention period**: Determine a length of time (days, hours, or minutes) for the system to retain the snapshot.

Note: When you set a retention period, you select a period that begins at the current time (retention is not calculated from the snapshot creation time).

11. Click **OK**.

Snapshot schedule details

You can view information about snapshot schedules on the **Schedules** page of the **Protection** tab from the **NetApp Element Management** extension point.

ID

System-generated ID for the schedule.

Type

Indicates the type of schedule. `Snapshot` is currently the only type supported.

Name

The name given to the schedule when it was created.

Frequency

The frequency at which the schedule is run. The frequency can be set in hours and minutes, days of the week, or days of the month.

Recurring

Indicates if the schedule is to run only once or at regular intervals.

Paused

Identifies whether or not the schedule has been manually paused.

Volume IDs

Displays the ID of the volume the schedule will use when the schedule is run.

Last Run

Displays the last time the schedule was run.

Last Run Status

Displays the outcome of the last schedule execution. Can be either `Success` or `Failure`.

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. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Schedules** sub-tab.
3. Select the check box for the snapshot schedule you want to edit.
4. Click **Actions**.
5. In the resulting menu, select **Edit**.
6. In the **Volume IDs CSV** field, modify the single volume ID or comma-separated list of volume IDs currently included in the snapshot operation.

7. Optional: To pause an active schedule or resume a paused schedule, select the **Manually Pause Schedule** check box.
8. Enter a different name for the schedule in the **New Schedule Name** field, if desired.
9. Change the current schedule type to one of the following, if desired:
    a. **Days of Week**: Select one of more days of the week and a time of day to create a snapshot.
    b. **Days of Month**: Select one of more days of the month and a time of day to create a snapshot.
    c. **Time Interval**: Select an interval for the schedule to run based on number of days, hours and minutes between snapshots.
10. Optional: Select **Recurrent Schedule** to repeat the snapshot schedule indefinitely.
11. Optional: Enter or modify the name for the snapshots defined by the schedule in the **New Snapshot Name** field.

    Note: If you leave the field blank, the system uses the time and date of the snapshot's creation as the name.
12. Optional: Select the **Include snapshots in replication when paired** check box to ensure that the snapshot are captured in replication when the parent volume is paired.
13. Optional: Select one of the following as the retention period for the snapshot:
    - **Keep forever**: Retains the snapshot on the system indefinitely.
    - **Set retention period**: Determine a length of time (days, hours, or minutes) for the system to retain the snapshot.

      Note: When you set a retention period, you select a period that begins at the current time (retention is not calculated from the snapshot creation time).
14. Click **OK**.

Copying a snapshot schedule

You can make a copy of a snapshot schedule and assign it to new volumes or use it for other purposes.

Steps

1. Select **NetApp Element Management > Protection**.

    Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.
2. Click the **Schedules** sub-tab.
3. Select the check box for the snapshot schedule you want to copy.
4. Click **Actions**.
5. In the resulting menu, click **Copy**.

    The **Copy Schedule** dialog box appears, populated with the current attributes of the schedule.
6. Optional: Enter a name and update attributes for the copy of the schedule.
7. Click **OK**.

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. Select **NetApp Element Management** > **Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Schedules** sub-tab.
3. Select the check box for the snapshot schedule you want to delete.
4. Click **Actions**.
5. In the resulting menu, click **Delete**.
6. Confirm the action.

Configuring cluster and volume pairing for real-time remote replication

For clusters running NetApp Element software, real-time replication enables the quick creation of remote copies of volume data. You can pair a storage cluster with up to four other storage clusters. You can replicate volume data synchronously or asynchronously from either cluster in a cluster pair for failover and failback scenarios. You must first pair two NetApp Element clusters and then pair volumes on each cluster to take advantage of real-time remote replication.

Before you begin

- You have added at least one cluster to the plug-in.
- 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 NetApp Element software versions on the clusters is no greater than one major version. If the difference is greater, one of the clusters must be upgraded to perform data replication.

Note: WAN Accelerator appliances have not been qualified by NetApp for use when replicating data. These appliances can interfere with compression and deduplication if deployed between two clusters that are replicating data. Be sure to fully qualify the effects of any WAN Accelerator appliance before you deploy it in a production environment

Steps

1. Pairing clusters on page 134
   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 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.

2. Pairing volumes on page 137
   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. You can pair the volume using known credentials for both clusters or by using a pairing key if cluster credentials are only available on the source cluster. If you know the credentials for both clusters, you can use a third option to create a replication target volume on the remote cluster to pair with the source cluster. After a volume pairing relationship is established, you must identify which volume is the replication target.
3. Validating volume replication on page 145
   After a volume is replicated, you should ensure that the source and target volumes are active. When in `Active` state, volumes are paired, data is being sent from the source to the target volume, and the data is in sync.
4. Deleting a volume relationship after replication on page 145
   After replication completes and you no longer need the volume pairing relationship, you can delete the volume relationship.
5. Managing volume relationships on page 145
   You can manage volume relationships in many ways, such as pausing replication, reversing volume pairing, changing the mode of replication, deleting a volume pair, or deleting a cluster pair.

Pairing clusters

You must pair two clusters as a first step to using real-time replication functionality. After you pair and connect two clusters, you can configure active volumes on one cluster to be continuously replicated to a second cluster, providing continuous data protection (CDP). 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

- You have Cluster Admin privileges to one or both clusters being paired.
- Ensure there is less than 2000 ms of round-trip latency between clusters.
- Ensure that the difference between NetApp Element software versions on the clusters is no greater than one major version.
- Ensure that all node IPs on paired clusters are routed to each other.

Note: Cluster pairing requires full connectivity between nodes on the management network. Replication requires connectivity between the individual nodes on the storage cluster network.

About this task

You can pair one NetApp Element cluster with up to four other clusters for replicating volumes. You can also pair clusters within the cluster group with each other.

Steps

1. Pairing clusters using known credentials on page 135
   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 clusters is required to use this method.
2. Pairing clusters with a pairing key on page 136
   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.

3. Validating the cluster pair connection on page 137
   After the cluster pairing has completed, you might want to verify the cluster pair connection to ensure replication success.

Pairing clusters using known credentials

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 clusters is required to use this method.

About this task

The Cluster Admin user name and password is used to authenticate cluster access before the clusters can be paired. If the MVIP is not known, or access to the cluster is not available, you can pair the cluster by generating a pairing key and use the key to pair the two clusters. For more detail, see plug-in documentation about pairing clusters with a pairing key.

Steps

1. From the vCenter that contains the local cluster, select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Cluster Pairs** sub-tab.
3. Click **Create Cluster Pairing**.
4. In the **Create Cluster Pairing** dialog box, select one of the following:
   - **Registered Cluster**: Select this method if the remote cluster of the pairing is controlled by the same instance of the NetApp Element Plug-in for vCenter Server.
   - **Credentialed Cluster**: Select this method if the remote cluster has known credentials that are outside of the NetApp Element Plug-in for vCenter Server configuration.
5. If you selected **Registered Cluster**, select a cluster from the list of available clusters and click **Pair**.
6. If you selected **Credentialed Cluster**, do the following:
   a. Enter the remote cluster MVIP address.
   b. Enter a cluster administrator user name.
   c. Enter a cluster administrator password.
   d. Click **Start Pairing**.
7. After the task completes and returns you to the **Cluster Pairs** page, verify that the cluster pair is connected.
8. Optional: On the remote cluster, select **NetApp Element Management > Protection > Cluster Pairs** or use the Element web UI to verify that the cluster pair is connected.

Related tasks

Pairing clusters with a pairing key on page 136

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.

Pairing clusters on page 134

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 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.

Pairing clusters with 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.

About this task

This procedure describes cluster pairing between two clusters using vCenter on the local and remote sites. For clusters not controlled by the vCenter Plug-in, you can alternately start or complete cluster pairing using the Element web UI. For detailed instructions on starting or completing cluster pairing from the Element web UI, see the user guide for NetApp Element software.

Steps

1. From the vCenter that contains the local cluster, select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Cluster Pairs** sub-tab.
3. Click **Create Cluster Pairing**.
4. In the **Create Cluster Pairing** dialog box, select **Inaccessible Cluster**.
5. 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.

6. Copy the cluster pairing key to your clipboard.
7. Click **Close**.
8. 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.

9. From the vCenter that contains the remote cluster, select **NetApp Element Management** > **Protection**

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

   Note: You can alternatively complete the pairing using the Element UI.

10. Click the **Cluster Pairs** sub-tab.
11. Click **Complete Cluster Pairing**.

    Note: Wait for the loading spinner to disappear before proceeding to the next step. If an unexpected error occurs during the pairing process, check for and manually delete any unconfigured cluster pairs on the local and remote cluster and perform the pairing again.

12. Paste the pairing key from the local cluster in the **Cluster Pairing Key** field.
13. Click **Pair Cluster**.
14. After the task completes and returns you to the **Cluster Pair** page, verify that the cluster pair is connected.
15. On the remote cluster, select **NetApp Element Management** > **Protection** or use the Element UI to verify that the cluster pair is connected.

Related tasks

Pairing clusters using known credentials on page 135
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 clusters is required to use this method.

Pairing clusters on page 134
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 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.

Validating the cluster pair connection

After the cluster pairing has completed, you might want to verify the cluster pair connection to ensure replication success.

Steps

1. On the local cluster, select **Data Protection** > **Cluster Pairs**.
2. In the Cluster Pairs window, verify that the cluster pair is connected.
3. Navigate back to the local cluster and the **Cluster Pairs** window and verify that the cluster pair is connected.

Pairing volumes

After you have established a connection between clusters in a cluster pair, you can pair a volume on one cluster with a volume on the other cluster in the pair. You can pair the volume using known credentials for both clusters or by using a pairing key if cluster credentials are only available on the source cluster. If you know the credentials for both clusters, you can use a third option to create a

replication target volume on the remote cluster to pair with the source cluster. After 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. Pairing volumes using known credentials on page 138
   You can pair a local volume with another volume on a remote cluster. Use this method if there is 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.
2. Creating target volumes and pairing them with local volumes on page 140
   You can pair two or more local volumes with associated target volumes on a remote cluster. This process creates a replication target volume on the remote cluster for each local source volume you select. Use this method if there is Cluster Admin access to both clusters on which volumes are to be paired and remote cluster is controlled by the plug-in. This method uses the volume ID of each volume on the remote cluster to initiate one or more connections.
3. Pairing volumes using a pairing key on page 142
   You can pair a local volume with another volume on a remote cluster using a pairing key. Use this method if there is 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.
4. Assigning a replication source and target to paired volumes on page 144
   If you did not assign a volume to be the replication target during volume pairing, configuration is not complete. You can use this procedure to 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 from a source volume to a remote target volume should the source volume become unavailable.

Pairing volumes using known credentials

You can pair a local volume with another volume on a remote cluster. Use this method if there is 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.

Before you begin

- You have Cluster Admin credentials for the remote cluster.
- Ensure that the clusters containing the volumes are paired.
- You know the remote Volume ID unless you intend to create a new volume during this process.
- If you intend for the local volume to be the source, ensure that the access mode of the volume is set to `Read/Write`.

Steps

1. From the vCenter that contains the local cluster, select **NetApp Element Management** > **Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.

3. From the **Active** view, select the check box for the volume you want to pair.
4. Click **Actions**.
5. In the resulting menu, click **Volume Pairing**
6. In the **Volume Pairing** dialog box, select one of the following:
   - **Volume Creation**: Select this method to create a replication target volume on the remote cluster. This method can only be used on remote clusters that are controlled by a NetApp Element Plug-in for vCenter Server.
   - **Volume Selection**: Select this method if the remote cluster for the target volume is controlled by a NetApp Element Plug-in for vCenter Server.
   - **Volume ID**: Select this method if the remote cluster for the target volume has known credentials that are outside of the NetApp Element Plug-in for vCenter Server configuration.
7. Select a Replication Mode from the drop-down list.
   - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
   - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
   - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.
8. If you selected **Volume Creation** as the pairing mode option, do the following:
   a. Select a paired cluster from the drop-down list.

      Note: This action populates the available accounts on the cluster to be selected in the next step.

   b. Select an account on the target cluster for the replication target volume.
   c. Enter a replication target volume name.

      Note: Volume size cannot be adjusted during this process.
9. If you selected **Volume Selection** as the pairing mode option, do the following:
   a. Select a paired cluster from the drop-down list.

      Note: This action populates the available volumes on the cluster to be selected in the next step.

   b. Optional: Click the **set remote volume to Replication Target** option if you want to set the remote volume as the target in the volume pairing. The local volume, if set to read/write, becomes the source in the pair.

      Attention: If you assign an existing volume as the replication target, the data on that volume will be overwritten. As a best practice, you should use a new volume as the replication target.

      Note: You can also assign replication source and target later in the pairing process from **Volumes > Actions > Edit**. You must assign a source and target to complete the pairing.

   c. Select a volume from the list of available volumes.
10. If you selected **Volume ID** as the pairing mode option, do the following:

a. Select a paired cluster from the drop-down list.

b. If the cluster is not registered with the plug-in, enter a cluster administrator user ID.

c. If the cluster is not registered with the plug-in, enter a cluster administrator password.

d. Enter a volume ID.

e. Click the **set remote volume to Replication Target** option if you want to set the remote volume as the target in the volume pairing. The local volume, if set to read/write, becomes the source in the pair.

Attention: If you assign an existing volume as the replication target, the data on that volume will be overwritten. As a best practice, you should use a new volume as the replication target.

Note: You can also assign replication source and target later in the pairing process from **Volumes** > **Actions** > **Edit**. You must assign a source and target to complete the pairing.

11. Click **Pair**.

Note: After you confirm the pairing, the two clusters begin the process of connecting the volumes. During the pairing process, you can see progress messages in the **Volume Status** column on the **Volume Pairs** page.

Attention: If you have not yet assigned a volume to be the replication target, the pairing configuration is not complete. The volume pair displays `PausedMisconfigured` until the volume pair source and target are assigned. You must assign a source and target to complete the volume pairing.

12. Select **Protection** > **Volume Pairs** on either cluster.

13. Verify the status of the volume pairing.

Related tasks

Assigning a replication source and target to paired volumes on page 144
If you did not assign a volume to be the replication target during volume pairing, configuration is not complete. You can use this procedure to 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 from a source volume to a remote target volume should the source volume become unavailable.

Related references

Volume pairing messages on page 149
Volume pairing warnings on page 149

Creating target volumes and pairing them with local volumes

You can pair two or more local volumes with associated target volumes on a remote cluster. This process creates a replication target volume on the remote cluster for each local source volume you select. Use this method if there is Cluster Admin access to both clusters on which volumes are to be paired and remote cluster is controlled by the plug-in. This method uses the volume ID of each volume on the remote cluster to initiate one or more connections.

Before you begin

- You have Cluster Admin credentials for the remote cluster.
- Ensure that the clusters containing the volumes are paired using the plug-in.

- Ensure that the remote cluster is controlled by the plug-in.
- Ensure that the access mode of each local volume is set to `Read/Write`.

Steps

1. From the vCenter that contains the local cluster, select **NetApp Element Management > Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for two or more volumes you want to pair.
4. Click **Actions**.
5. In the resulting menu, click **Volume Pairing**.
6. In the **Multiple Volume Pairing with Target Volume Creation** dialog box, select a **Replication Mode** from the drop-down list.
   - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
   - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
   - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.
7. Select a paired cluster from the drop-down list.
8. Select an account on the target cluster for the replication target volume.
9. Optional: Type a prefix or suffix for the new volume names on the target cluster.

   Note: A sample volume name with the modified name is shown.

10. Click **Create Pairs**.

    Note: After you confirm the pairing, the two clusters begin the process of connecting the volumes. During the pairing process, you can see progress messages in the **Volume Status** column on the **Volume Pairs** page. After the process completes, new target volumes are created and connected on the remote cluster.

11. Select **Protection > Volume Pairs** on either cluster.
12. Verify the status of the volume pairing.

Related references

Volume pairing messages on page 149
Volume pairing warnings on page 149

Pairing volumes using a pairing key

You can pair a local volume with another volume on a remote cluster using a pairing key. Use this method if there is 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.

Before you begin

Ensure that the clusters containing the volumes are paired.

Best Practices: Set the source volume to `Read/Write` and the target volume to `Replication Target`. The target volume should contain no data and have 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 may be greater or equal in size to the source volume, but it cannot be smaller.

About this task

This procedure describes volume pairing between two volumes using vCenter on the local and remote sites. For volumes not controlled by the vCenter Plug-in, you can alternately start or complete volume pairing using the Element web UI. For detailed instructions on starting or completing volume pairing from the Element web UI, see the user guide for NetApp Element software.

Note: The volume pairing key contains an encrypted version of the volume information and may contain sensitive information. Share this key only in a secure manner.

Steps

1. From the vCenter that contains the local cluster, select **NetApp Element Management** > **Management**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volumes** sub-tab.
3. From the **Active** view, select the check box for the volume you want to pair.
4. Click **Actions**.
5. In the resulting menu, click **Volume Pairing**
6. In the **Volume Pairing** dialog box, select **Inaccessible Cluster**.
7. Select a **Replication Mode** from the list.

   - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
   - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
   - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.

8. 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.

9. Copy the pairing key to your clipboard.
10. Click **Close**.
11. 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 stored 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.

12. From the vCenter that contains the remote cluster, select **NetApp Element Management** > **Management**.

    Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

13. Click the **Volumes** sub-tab.
14. From the **Active** view, select the check box for the volume you want to pair.
15. Click **Actions**.
16. In the resulting menu, click **Volume Pairing**.
17. In the **Volume Pairing** dialog box, select **Complete Cluster Pairing**.
18. Paste the pairing key from the other cluster into the **Pairing Key** box.
19. Click **Complete Pairing**.

    Note: After you confirm the pairing, the two clusters begin the process of connecting the volumes. During the pairing process, you can see progress messages in the **Volume Status** column of the **Volume Pairs** page. If an unexpected error occurs during the pairing process, check for and manually delete any unconfigured cluster pairs on the local and remote cluster and perform the pairing again.

    Attention: If you have not yet assigned a volume to be the replication target, the pairing configuration is not complete. The volume pair displays `PausedMisconfigured` until the volume pair source and target are assigned. You must assign a source and target to complete the volume pairing.

20. Select **Protection** > **Volume Pairs** on either cluster.
21. Verify the status of the volume pairing.

    Note: Volumes that are paired using a pairing key appear after the pairing process has been completed at the remote location.

Related tasks

Assigning a replication source and target to paired volumes on page 144

If you did not assign a volume to be the replication target during volume pairing, configuration is not complete. You can use this procedure to 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 from a source volume to a remote target volume should the source volume become unavailable.

Related references

Volume pairing messages on page 149

Volume pairing warnings on page 149

Assigning a replication source and target to paired volumes

If you did not assign a volume to be the replication target during volume pairing, configuration is not complete. You can use this procedure to 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 from 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.

About this task

This procedure describes assigning source and replication volumes between two clusters using vCenter on the local and remote sites. For volumes not controlled by the vCenter Plug-in, you can alternately assign a source or replication volume using the Element web UI. For detailed instructions on deleting a cluster pair end from the Element web UI, see the user guide for NetApp Element software.

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

Best Practices: The target volume should contain no data and have the exact characteristics of the source volume, such as size, 512e setting, and QoS configuration. The target volume may be greater or equal in size to the source volume, but it cannot be smaller.

Steps

1. From the **NetApp Element Management** extension point, select the cluster containing the paired volume that you want to use as the replication source.
2. Select **NetApp Element Management** > **Management**.
3. Click the **Volumes** sub-tab.
4. From the **Active** view, select the check box for the volume you want to edit.
5. Click **Actions**.
6. In the resulting menu, select **Edit**.
7. 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 `PausedMisconfigured` until a new replication target is assigned. Changing access pauses volume replication and causes the transmission of data to cease. Be sure that you have coordinated these changes at both sites.

8. Click **OK**.
9. From the **NetApp Element Management** extension point, select the cluster containing the paired volume that you want to use as the replication target.
10. Select **NetApp Element Management** > **Management**.
11. Click the **Volumes** sub-tab.
12. From the **Active** view, select the check box for the volume you want to edit.
13. Click **Actions**.

14. In the resulting menu, select **Edit**.
15. 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. As a best practice, you should use a new volume as the replication target.

16. Click **OK**.

Related references

Volume pairing messages on page 149
Volume pairing warnings on page 149

Validating volume replication

After a volume is replicated, you should ensure that the source and target volumes are active. When in `Active` state, volumes are paired, data is being sent from the source to the target volume, and the data is in sync.

Steps

1. From either cluster in the pairing, select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volume Pairs** sub-tab.
3. Verify that the volume status is `Active`.

Deleting a volume relationship after replication

After replication completes and you no longer need the volume pairing relationship, you can delete the volume relationship.

Step

1. Complete the procedure that describes how to delete a volume pairing relationship. *Deleting volume pairs*

Managing volume relationships

You can manage volume relationships in many ways, such as pausing replication, reversing volume pairing, changing the mode of replication, deleting a volume pair, or deleting a cluster pair.

Related tasks

Pausing replication on page 146
Changing the mode of replication on page 146
Deleting volume pairs on page 147

Related references

Volume pair details on page 148

Pausing replication

You can edit volume pair properties to manually pause replication.

Steps

1. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volume Pairs** sub-tab.
3. Select the check box for the volume pair you want to edit.
4. Click **Actions**.
5. In the resulting menu, select **Edit**.
6. In the **Edit Volume Pair** pane, manually pause or start the replication process:

   Attention: Pausing or resuming volume replication manually will cause the transmission of data to cease or resume. Be sure that you have coordinated these changes at both sites.

7. Click **Save Changes**.

Changing the mode of replication

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

Steps

1. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volume Pairs** sub-tab.
3. Select the check box for the volume pair you want to edit.
4. Click **Actions**.
5. In the resulting menu, select **Edit**.
6. In the **Edit Volume Pair** pane, select a new replication mode:

   Attention: Changing the mode of replication will cause the mode to change immediately. Be sure that you have coordinated these changes at both sites.

   - **Real-time (Synchronous)**: Writes are acknowledged to the client after they are committed on both the source and target clusters.
   - **Real-time (Asynchronous)**: Writes are acknowledged to the client after they are committed on the source cluster.
   - **Snapshots Only**: Only snapshots created on the source cluster are replicated. Active writes from the source volume are not replicated.

7. Click **Save Changes**.

Deleting volume pairs

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

About this task

This procedure describes deleting a volume pairing relationship between two volumes using vCenter on the local and remote sites. For volumes not controlled by the vCenter Plug-in, you can alternately delete a volume pair end using the Element web UI. For detailed instructions on deleting a volume pair end from the Element web UI, see the user guide for NetApp Element software.

Steps

1. Select **NetApp Element Management > Protection**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Volume Pairs** sub-tab.
3. Select the check box for one or more of the volume pairs you want to delete.
4. Click **Actions**.
5. In the resulting menu, select **Delete**.
6. In the **Delete Volume Pair** dialog box, confirm the details of each volume pair.

   Note: For clusters that are not managed by the plug-in, this action deletes only the volume pair end on the local cluster. You need to manually delete the volume pair end from the remote cluster to fully remove the pairing relationship.

7. (Optional for clusters managed by plug-in) Select the check box **Change Replication Target Access to** and select a new access mode for the replication target volume. This new access mode will be applied after the volume pairing relationship has been removed.
8. Click **Yes**.

Deleting a cluster pair

You can delete a cluster pairing relationship between two clusters using vCenter on the local and remote sites. To completely remove a cluster pairing relationship, you must remove cluster pair ends from both the local and remote clusters.

About this task

You can use the vCenter Plug-in to delete a cluster pair end. For clusters not controlled by the vCenter Plug-in, you can alternately delete a cluster pair end using the Element web UI. For detailed instructions on deleting a cluster pair end using the Element web UI, see the user guide for NetApp Element software.

Steps

1. From either cluster, select **NetApp Element Management > Protection**.
2. Click the **Cluster Pairs** sub-tab.
3. Select the check box for the cluster pair you want to delete.
4. Click **Actions**.
5. In the resulting menu, click **Delete**.

6. Confirm the action.

   Note: This action deletes only the cluster pair end on the local cluster. You need to manually delete the cluster pair end from the remote cluster to fully remove the pairing relationship.

7. Perform the steps again from the remote cluster in the cluster pairing.

Volume pair details

You can view the information for volumes that have been paired or are in the process of being paired on the **Volume Pairs** page of the **Protection** tab from the **NetApp Element Management** configuration point.

The system displays pairing and progress messages in the **Volume Status** column.

Local Volume ID

System-generated ID for the volume.

Local 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 Replication status

Replication status of the volume.

Snapshot Replication status

Replication status of the snapshot volume.

Mode

Indicates the client write replication method. Possible values:

- `Asynchronous`
- `Snapshot-Only`
- `Synchronous`

Local Volume Access

The access mode of the local half of the volume pairing.

Direction

Indicates the direction of the volume data:

- The icon  indicates data is being written to a target outside the cluster.
- The icon  indicates data is being written to the local volume from an outside source.

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.

Related references

Volume pairing messages on page 149
Volume pairing warnings on page 149

Volume pairing messages

You can view messages during the initial pairing process on the **Volume Pairs** page of the **Protection** tab from the **NetApp Element Management** configuration point. These messages are displayed in the **Volume Status** column and can display on both source and target ends of the pairing.

PausedDisconnected

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

ResumingConnected*

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

ResumingRRSync*

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

ResumingLocalSync*

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

ResumingDataTransfer*

Data transfer has been 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.

*This process is driven by the target volume and might not display on the source volume.

Related references

Volume pairing warnings on page 149

Volume pairing warnings

You can view warning messages after you pair volumes on the **Volume Pairs** page of the **Protection** tab from the **NetApp Element Management** configuration point. These messages are displayed in the **Volume Status** column and can display on both source and target ends of the pairing.

These messages can display on both source and target ends of the pairing unless otherwise indicated.

PausedClusterFull

Because the target cluster is full, source replication and bulk data transfer cannot proceed. The message displays on the source end of the pair only.

PausedExceededMaxSnapshotCount

The target volume already has the maximum number of snapshots and cannot replicate additional snapshots.

PausedManual

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

PausedManualRemote

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

PausedMisconfigured

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

PausedQoS

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

PausedSlowLink

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

PausedVolumeSizeMismatch

Target volume is 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.

StoppedMisconfigured

A permanent configuration error has been detected. The remote volume has been purged or unpaired. No corrective action is possible; a new pairing must be established.

Related references

Volume pairing messages on page 149

Cluster management

From the **Cluster** tab in the **NetApp Element Management** extension point, you can view and change cluster-wide settings and perform cluster-specific tasks.

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.

A storage node contains the following types of drives:

Volume metadata drives

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

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

Block drives

These drives store the compressed, deduplicated data blocks for server application volumes. Block drives make up a majority of the storage capacity of the system. The majority of read requests for data already stored on the cluster, as well as requests to write data, occur on the block drives. The total block drive capacity in the system determines the maximum amount of data that can be stored, taking into account the effects of compression, thin provisioning, and deduplication.

Adding available drives to a cluster

You can add drives to a cluster using the NetApp Element Management extension point. 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 before each drive can participate in the cluster.

About this task

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

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

Steps

1. Select **NetApp Element Management** > **Cluster**

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Drives** sub-tab, select **Available** from the drop-down list to view the list of available drives.
3. Add drives as follows:

a. Select the check box for each drive you want to add.

b. Click **Add Drives**.

4. Review the details of the drives you are intending to add and confirm the action.

Drive details

You can view a list of the active drives in the cluster using the **Active** view on the **Drives** page of the **Cluster** tab from the **NetApp Element Management** extension point. You can change the view by selecting available options using the drop-down filter.

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 cluster is created. The following headings are shown in the list of active drives:

Drive ID

Sequential number assigned to the drive.

Drive State

The status of the drive.

Possible values are as follows:

- `Active`: A drive that is in use by the cluster.
- `Available`: An available drive that can be added to the cluster.
- `Removing`: A drive is in the process of being removed. Any data previously on the drive is being migrated to other drives in the cluster.
- `Erasing`: A drive is in the process of being secure erased. Any data on that drive is permanently removed
- `Failed`: A drive that has failed. Any data that was previously on the drive has been migrated to other drives in the cluster. The system puts a drive in a failed state if the self-diagnostics of the drive tells the node it has failed or if communications to the drive stop for 5.5 minutes or longer.

Node ID

Assigned node number when the node is added to the cluster.

Node Name

Name of the node where the drive resides.

Slot

Slot number where the drive is physically located.

Capacity (GB)

Size of the drive in GB.

Serial

Serial number of the SSD.

Wear Remaining

Wear-level indicator.

Type

Drive type can be block or metadata.

Removing a drive

You can remove a drive from a cluster using the NetApp Element Management extension point. You might do this when reducing cluster capacity or preparing to replace drives nearing the end of their service life. Removing a drive takes the drive offline. Any data on the drive is removed and migrated to other drives in the cluster before the drive is removed from the cluster. The data migration to other active drives in the system can take a few minutes to an hour depending on capacity utilization and active I/O on the cluster.

About this task

When you remove a drive in a `Failed` state, the drive is not returned to `Available` or `Active` states. Instead, the drive is unavailable for use in the cluster.

Steps

1. Select **NetApp Element Management > Cluster**.
2. If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.
3. Select **All** from the drop-down list to view the complete list of drives.
4. Remove drives as follows:
   a. Select the check box for each drive you want to remove.
   b. Click **Remove Drives**.
5. Confirm the action.

   Note: If there is not enough capacity to remove active drives before removing a node, an error message appears when you confirm the drive removal.

Nodes

Nodes are the hardware that is grouped into a cluster to be accessed as block storage or compute resources.

NetApp HCI storage nodes

NetApp HCI storage nodes are hardware that provide the storage resources for a NetApp HCI system. Drives in the node contain block and metadata space for data storage and data management. Each node contains a factory image of NetApp Element software. NetApp HCI storage nodes can be managed using the NetApp Element Management extension point.

NetApp HCI compute nodes

NetApp HCI compute nodes are node hardware that provide the resources, such as CPU, memory, and networking, that are needed for virtualization in the NetApp HCI installation. Because each server runs VMware ESXi, NetApp HCI compute node management (adding or removing hosts) must be done outside of the plug-in within the Hosts and Clusters menu in vSphere.

SolidFire Storage nodes

A SolidFire all-flash array storage node is a collection of drives that communicate with each other through the CIPI Bond10G network interface. Drives in the node contain block and metadata space

for data storage and data management. You can create a cluster with new storage nodes, or add storage nodes to an existing cluster to increase storage capacity and performance.

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 add or remove nodes from the cluster at any time without interrupting service.

SolidFire 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.

Fibre Channel nodes are added in pairs, and operate in active-active mode (all Fibre Channel nodes actively process traffic for the cluster). At least two Fibre Channel nodes are required for Fibre Channel connectivity in a SolidFire cluster. Clusters running NetApp Element software 9.0 and later support up to four Fibre Channel nodes.

Adding a node to a cluster

You can add storage nodes when a cluster is created or when more storage is needed.

Before you begin

- The node you are adding has been set up, powered on, and configured.
- Both the major or minor version numbers of the software on each node in a cluster must match for the software to be compatible. For example, Element 9.0 is not compatible with version 9.1.

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

About this task

Nodes require initial configuration when they are first powered on. When the node has been set up and configured, it registers itself on the cluster identified when the node was configured and appears in the list of pending nodes on the **Cluster > Nodes** page of the NetApp Element Management extension point.

You can add nodes of smaller or larger capacities to an existing cluster.

The procedure is the same for adding Fibre Channel nodes or storage nodes that are running NetApp Element software.

Steps

1. Select **NetApp Element Management > Cluster**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Nodes** sub-tab.
3. Select **Pending** from the drop-down list to view the list of nodes.
4. To add one or more nodes, perform the following steps:
   a. Select the check box for each node you want to add.
   b. Click **Add Node**.
5. Review the details of the nodes you are intending to add and confirm the action.

   When the action is complete, the node appears in the list of active nodes for the cluster.

Node details

You can view a list of the nodes in the cluster on the **Nodes** page of the **Cluster** tab from the **NetApp Element Management** extension point. You must select **Active** view to see the list of active nodes. You can change the view by selecting **Pending**, **PendingActive**, and **All** options using the drop-down filter.

Node ID

System-generated ID for the node.

Node Name

The system-generated or administrator-assigned node name.

Node State

The status of the node.

Possible values:

- **Active**: The node is an active member of a cluster and cannot be added to another cluster.
- **Pending**: The node is pending for a specific named cluster and can be added.
- **PendingActive**: The node is currently being returned to the factory software image, and is not yet an active member of a cluster. When complete, it will transition to the **Active** state.

Available 4k IOPS

Displays the IOPS configured for the node.

Node Role

Identifies what role the node has in the cluster. The node can be a cluster master, ensemble member, or a Fibre Channel node.

Node Type

Displays the model type of the node.

Active Drives

Number of active drives in the node.

Management IP

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

Storage IP

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

The virtual ID for the storage local area network.

Version

Version of NetApp Element software running on each node.

Removing nodes from a cluster

You can remove nodes from a cluster without service interruption when their storage is no longer needed or they require maintenance.

Before you begin

You have removed all the drives in the node from the cluster. You cannot remove a node until the *RemoveDrives* process has completed and all data has been migrated away from the node.

About this task

At least two Fibre Channel nodes are required for Fibre Channel connectivity in a NetApp Element cluster. If only one Fibre Channel node is connected, the system triggers alerts in the Event Log until you add another Fibre Channel node to the cluster, even though all Fibre Channel network traffic continues to operate with only one Fibre Channel node.

Steps

1. Select **NetApp Element Management > Cluster**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Nodes** sub-tab.
3. To remove one or more nodes, perform the following steps:
   a. From **Active** view, select the check box for each node you want to remove.
   b. Click **Actions**.
   c. Select **Remove**.
4. Confirm the action.

   Any nodes removed from a cluster appear in the list of **Pending** nodes.

Related tasks

Removing a drive on page 153

Restarting cluster nodes

You can restart one or more active nodes in a cluster using the NetApp Element Management extension point.

Before you begin

You have stopped I/O and disconnected all iSCSI sessions if you are restarting more than one node simultaneously.

About this task

To restart the cluster, you can select all cluster nodes and perform a restart.

Attention: This method restarts all networking services on a node, causing temporary loss of networking connectivity.

Steps

1. Select **NetApp Element Management > Cluster**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Nodes** sub-tab.
3. From **Active** view, select the check box for each node you want to restart.
4. Click **Actions**.
5. Select **Restart**.
6. Confirm the action.

Shutting down cluster nodes

You can shut down one or more active nodes in a cluster using the NetApp Element Management extension point. To shut down the cluster, you can select all cluster nodes and perform a simultaneous shutdown.

Before you begin

You have stopped I/O and disconnected all iSCSI sessions if you are shutting down more than one node simultaneously.

Steps

1. Select **NetApp Element Management > Cluster**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Nodes** sub-tab.
3. From **Active** view, select the check box for each node you want to shut down.
4. Click **Actions**.
5. Select **Shutdown**.
6. Confirm the action.

Note: If a node has been down longer than 5.5 minutes under any type of shutdown condition, the NetApp Element software determines that the node is not coming back to join the cluster. Double Helix data protection begins the task of writing single replicated blocks to another node to replicate the data. Depending on the length of time a node is shut down, its drives might need to be added back to the cluster after the node is brought back online.

Related tasks

Restarting cluster nodes on page 157

Related information

Powering off and powering on a NetApp HCI system

VLAN management

Virtual networking in NetApp Element storage allows traffic between multiple clients that are on separate logical networks to be connected to one NetApp Element cluster. To implement virtual networking, Element uses VLAN backing technology.

The NetApp Element Plug-in for vCenter enables you to manage VLANs for the selected cluster. You can create, view, edit, and delete VLANs. VLAN management options are available only from the NetApp Element Management extension point.

Creating a VLAN

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

Before you begin

- ESXi hosts have a single iSCSI software adapter.
- Hosts or switches are configured for the VLAN.
- You have identified the block of IP addresses that will be assigned to the virtual networks on the cluster nodes.
- You have identified a storage network IP (SVIP) address that will be used as an endpoint for all NetApp Element storage traffic.

Attention: The following criteria should be considered for this configuration:

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

About this task

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

Steps

1. Select **NetApp Element Management > Cluster**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Network** sub-tab.
3. Click **Create VLAN**.
4. In the **Create VLAN** dialog box, enter a name for the VLAN.
5. Enter an integer for the VLAN tag.
6. Enter the Storage Virtual IP (SVIP) address for the storage cluster.
7. Adjust the netmask, as needed.

   The default is

   `255.255.255.0`

8. Optional: Enter a description for the VLAN.
9. Optional: Select the **Enable Virtual Routing and Forwarding** check box.

   Note: Virtual routing and forwarding (VRF) allows multiple instances of a routing table to exist in a router and work simultaneously. This functionality is available for storage networks only.

   a. Enter an IP address of a gateway of the virtual network.

10. Select the hosts that you want to include in the VLAN.

    Note: If you are using vCenter Linked Mode, only hosts available to the vCenter Server to which the cluster is assigned are available to select.

11. Configure the IP address blocks for the storage nodes as follows:

    Note: A minimum of one IP address block must be created.

    a. Click **Create Block**.

    b. Enter the starting address for the IP range.

    c. Enter the number of IP addresses to include in the address block.

       Note: The total number of IP addresses must match the number of nodes in the storage cluster.

    d. Click outside the entry to accept the values.

12. Click **OK** to create the VLAN.

Virtual network details

You can view network information for VLANs on the **Network** page of the **Cluster** tab from the **NetApp Element Management** extension point.

ID

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

VLAN 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.

IPs Used

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

Editing a virtual network

You can change VLAN attributes, such as VLAN name, netmask, and size of the IP address blocks.

About this task

The VLAN Tag and SVIP cannot be modified for a VLAN. The gateway attribute can only be modified for VRF VLANs. If any iSCSI, remote replication, or other network sessions exist, the modification might fail.

Steps

1. Select **NetApp Element Management > Cluster**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Network** sub-tab
3. Select the check box for the VLAN you want to edit.
4. Click **Actions**.
5. In the resulting menu, click **Edit**.
6. In the resulting menu, enter the new attributes for the VLAN.
7. Click **Create Block** to add a non-continuous block of IP addresses for the virtual network.
8. Click **OK**.

Deleting a virtual network

You can permanently delete a VLAN object and its block of IPs. Address blocks that were assigned to the VLAN are disassociated with the virtual network and can be reassigned to another virtual network.

Steps

1. Select **NetApp Element Management > Cluster**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Network** sub-tab.
3. Select the check box for the VLAN you want to delete.
4. Click **Actions**.
5. In the resulting menu, click **Delete**.
6. Confirm the action.

Virtual volumes

You can view information and perform tasks for virtual volumes and their associated storage containers, protocol endpoints, bindings, and hosts from the NetApp Element Management extension point. You must enable Virtual Volume functionality for it to be available on the cluster.

Configuring vSphere Virtual Volumes (VVols) functionality

You must perform initial configuration steps to use virtual volumes (VVols) in the NetApp Element Plug-in for vCenter Server.

Before you begin

The NetApp Element cluster must be connected to an ESXi 6.0 or later environment that is compatible with VVols.

Steps

1. Enable the virtual volumes feature on the NetApp Element cluster using the NetApp Element Configuration extension point.
2. Register the VASA provider with vCenter.
3. Create a storage container and associated VVol datastore using the NetApp Element Management extension point.
4. If you elected not to create a datastore when you created a storage container, create a VVol datastore in vCenter and associate it with the storage container.

Related tasks

Enabling virtual volumes on page 53
Registering the NetApp Element VASA Provider on page 161
Creating a storage container on page 168
Creating a VVol datastore on page 162

Registering the NetApp Element VASA Provider

You must register the NetApp Element VASA Provider with vCenter so that vCenter is aware of VVol functionality on the cluster. Registering the VASA provider with vCenter is a one-time configuration task.

Before you begin

- You have enabled the VVols feature on the cluster.
- You have vCenter 6.x.
- You have ESXi 6.x.

Attention: Do not register a NetApp Element VASA provider to more than one vCenter instance. The NetApp Element VASA provider can only be registered to a single vCenter due to limitations with how vCenter handles SSL. A single vCenter can have multiple NetApp Element clusters, but a cluster cannot be shared between two instances of vCenter.

About this task

This procedure describes the steps available in version 6.7 of vSphere. Your vSphere user interface may differ slightly from what is described depending on the version of vSphere installed. For additional help, see VMware vCenter documentation.

Steps

1. From vCenter Server, select **Hosts and Clusters**.
2. Select a vCenter server to register the NetApp Element VASA Provider.
3. Select **Configure** > **Storage Providers**.
4. From **Storage Providers**, click the add icon.

   The **New Storage Provider** dialog box appears.
5. Enter the following information:
   - VASA Provider name.
   - VASA Provider URL.

     Note: The VASA Provider URL is provided to you when you enable VVols in the vCenter Plug-in. It is also available in **NetApp Element Configuration** > **Clusters** if you click **Actions** for the cluster you are enabling and click **Details**.
   - Administrative account user name for the NetApp Element cluster.
   - Administrative account password for the NetApp Element cluster.
   - Click **OK** to add the VASA Provider.
6. Approve the thumbprint of the SSL cert when prompted.

   The NetApp Element VASA Provider should now be registered with a status of `Connected`.

   Note: Refresh the storage provider, if necessary, to show the current status of the provider after registering the provider for the first time. You can also verify that the provider is enabled in **NetApp Element Configuration** > **Clusters**. Click **Actions** for the cluster you are enabling and click **Details**.

Creating a VVol datastore

You must create a virtual volume datastore that represents the storage container on the NetApp Element cluster in vCenter. You can create a VVol datastore using the Create Storage Container wizard or by using this process. You must create at least one VVol datastore to begin provisioning VVol-backed virtual machines.

Before you begin

- A cluster running NetApp Element software with VVols functionality enabled.
- An existing storage container in the virtual environment.
- The VASA provider must be registered with vCenter.

Note: You might need to rescan NetApp Element storage in vCenter to discover storage containers.

Steps

1. From the **Navigator** view in vCenter, right-click a storage cluster and select **Storage** > **Datastores** > **New Datastore**.

2. In the **New Datastore** dialog box, select **VVol** as the type of datastore to create.
3. Provide a name for the datastore in the **Datastore name** field.
4. Select the NetApp Element storage container from the **Backing Storage Container** list.

   Note: You do not need to manually create protocol endpoint (PE) LUNs. They are automatically mapped to the ESXi hosts when the datastore is created.

5. Select the hosts that require access to the datastore.
6. Click **Next**.
7. Review the configurations and click **Finish** to create the VVol datastore.

Related tasks

Enabling virtual volumes on page 53
Registering the NetApp Element VASA Provider on page 161
Creating a storage container on page 168

Viewing virtual volume details

You can review general information for all active virtual volumes on the cluster in the NetApp Element Management extension point. You can also view details specific to each virtual volume, including efficiency, performance, and QoS as well as associated snapshots, parent virtual machine, bindings, and task status.

Before you begin

- You have completed VVol enablement and VASA provider registration steps.
- You have created at least one virtual volume.
- You have powered on VMs so virtual volume details are available to view.

Steps

1. Select **NetApp Element Management > VVols**.

   If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. From the **Virtual Volumes** tab, you can search for a specific virtual volume.
3. Select the check box for the virtual volume you want to review.
4. Click **Actions**.
5. In the resulting menu, select **Details**.

Related tasks

Configuring vSphere Virtual Volumes (VVols) functionality on page 161

Virtual volume details

You can view virtual volume information for all active virtual volumes on the cluster on the **Virtual Volumes** page of the **VVols** tab from the **NetApp Element Management** extension point.

Virtual Machine ID

The UUID of the virtual machine.

Name

The name assigned to the virtual volume.

Type

The virtual volume type: Config, Data, Memory, Swap, or Other.

Container

The UUID of the storage container that owns the virtual volume.

Volume ID

The ID of the underlying volume.

Virtual Volume ID

The UUID of the virtual volume.

Individual virtual volume details

You can view virtual volume information for an individual virtual volume when you select it and view its details on the **Virtual Volumes** page of the **VVols** tab from the **NetApp Element Management** extension point.

The following details are in the **General Details** section:

Name

The name assigned to the virtual volume.

Volume ID

The ID of the underlying volume.

Virtual Volume ID

The UUID of the virtual volume.

Virtual Volume Type

The virtual volume type: Config, Data, Memory, Swap, or Other.

Status

The status of VVol task.

Storage Container

The UUID of the storage container that owns the virtual volume.

Size

Size of the volume in GB or GiB.

Access

The read/write permissions assigned to the virtual volume.

The following details are in the **Efficiency** section:

Compression

The compression efficiency score for the volume.

Deduplication

The deduplication efficiency score for the volume.

Thin Provisioning

The thin provisioning efficiency score for the volume.

Last Updated

The date and time of the last efficiency score.

The following details are in the **Performance** section:

Account ID

The unique account ID of the associated account.

Actual IOPS

Current actual IOPS to the volume in the last 500 milliseconds.

Async Delay

The length of time since the volume was last synced with the remote cluster.

Average IOP Size

Average size in bytes of recent I/O to the volume in the last 500 milliseconds.

Burst IOPS Size

The total number of IOP credits available to the user. When volumes are not using up to the Max IOPS, credits are accrued.

Client Queue Depth

The number of outstanding read and write operations to the volume.

Last Updated

The date and time of the last performance update.

Latency USec

The average time, in microseconds, to complete operations to the volume in the last 500 milliseconds. A "0" (zero) value means there is no I/O to the volume.

Non-Zero Blocks

Total number of 4KiB blocks with data after the last garbage collection operation has completed.

Performance Utilization

The percentage of cluster IOPS being consumed. For example, a 250K IOP cluster running at 100K IOPS would show 40% consumption.

Read Bytes

The total cumulative bytes read from the volume since the creation of the volume.

Read Latency USec

The average time, in microseconds, to complete read operations to the volume in the last 500 milliseconds.

Read Operations

The total read operations to the volume since the creation of the volume.

Thin Provisioning

The thin provisioning efficiency score for the volume.

Throttle

A floating value between 0 and 1 that represents how much the system is throttling clients below their maxIOPS because of re-replication of data, transient errors and snapshots taken.

Total Latency USec

The time, in microseconds, to complete read and write operations to a volume.

Unaligned Reads

For 512e volumes, the number of read operations that were not on a 4k sector boundary. High numbers of unaligned reads may indicate improper partition alignment.

Unaligned Writes

For 512e volumes, the number of write operations that were not on a 4k sector boundary. High numbers of unaligned writes may indicate improper partition alignment.

Volume ID

The system-generated ID for the volume.

Volume Size

Size of the volume in bytes.

Volume Utilization

A percentage value that describes how much the client is using the volume.

Possible Values:

- 0 = Client is not using the volume
- 100 = Client is using the max
- >100 = Client is using the burst

Write Bytes

The total cumulative bytes written to the volume since the creation of the volume.

Write Latency USec

The average time, in microseconds, to complete write operations to a volume in the last 500 milliseconds.

Write Operations

The total cumulative write operations to the volume since the creation of the volume.

Zero Blocks

Total number of 4KiB blocks without data after the last round of garbage collection operation has completed.

Used Capacity

Percentage of used capacity.

The following details are in the **Quality of Service** section:

I/O Size

The size of the IOPS in KB.

Min IOPS

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

Max IOPS

The maximum number of sustained IOPS that the 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.

Max Bandwidth

The maximum bandwidth permitted by the system to process larger block sizes.

The following details are in the **Virtual Machine** section:

Virtual Machine ID

The UUID of the virtual machine.

VM Name

The name of the virtual machine.

Guest OS Type

Operating system associated with the virtual volume.

Virtual Volumes

List of virtual volumes UUIDs and VM names.

These headings are in the **Bindings** section:

Host

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 cluster.

PE Type

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

The following details are in the **Tasks** section:

Operation

The type of operation the task is performing.

Values:

- `unknown`: The task operation is unknown.
- `prepare`: The task is preparing a virtual volume.
- `snapshot`: The task is creating a snapshot of a virtual volume.
- `rollback`: The task is rolling back a virtual volume to a snapshot.
- `clone`: The task is creating a clone of the virtual volume.
- `fastClone`: The task is creating a fast clone of a virtual volume.
- `copyDiffs`: The task is copying differing blocks to a virtual volume.

Status

The current status of the virtual volume task:

Values:

- `Error`: The task has failed and returned an error.
- `Queued`: The task is waiting to be run.
- `Running`: The task is currently running.
- `Success`: The task has completed successfully.

Task ID

The unique ID of the task.

Storage containers

Storage containers are logical constructs that map to NetApp Element accounts and are used for reporting and resource allocation. They pool raw storage capacity or aggregate storage capabilities that the storage system can provide to virtual volumes. A VVol datastore that is created in vSphere is mapped to an individual storage container. A single storage container has all available resources from the NetApp Element cluster by default. If more granular governance for multi-tenancy is required, multiple storage containers can be created.

Storage containers function like traditional accounts and can contain both virtual volumes and traditional volumes. A maximum of four storage containers per cluster is supported. A minimum of one storage container is required to use VVols functionality. On the **VVols > Storage Containers** page of the NetApp Element Management extension point, you can create, delete, and view details about storage containers. You can discover storage containers in vCenter during VVols creation.

Creating a storage container

You can create storage containers from the **VVols** tab in the **NetApp Element Management** extension point. You must create at least one storage container to begin provisioning VVol-backed virtual machines.

Before you begin

- You have enabled VVols functionality for the cluster.
- You have registered the NetApp Element VASA Provider for virtual volumes with vCenter.

Steps

1. Select **NetApp Element Management > VVols**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Storage Containers** sub-tab.
3. Click **Create Storage Container**.
4. Enter storage container information in the **Create a New Storage Container** dialog box:
   a. Enter a name for the storage container.

      Tip: Use descriptive naming best practices. This is especially important if multiple clusters or vCenter Servers are used in your environment.

   b. Configure initiator and target secrets for CHAP.

      Best Practices: Leave the **CHAP Settings** fields blank to automatically generate secrets.

   c. Optional: Enter a name for the datastore. The **Create a datastore** check box is selected by default.

      Note: A VVol datastore is required to use the storage container in vSphere. If you choose not to create a datastore, you must create one later using the vSphere New Datastore wizard.

   d. Select one or more hosts for the datastore.

Note: If you are using vCenter Linked Mode, only hosts available to the vCenter Server to which the cluster is assigned are available to select.

e. Click **OK**.

5. Verify that the new storage container appears in the list in the **Storage Containers** sub-tab.

   Note: Because a NetApp Element account ID is created automatically and assigned to the storage container, it is not necessary to manually create an account.

6. Optional: If you selected **Create a datastore**, verify that the associated datastore has also been created on the selected host in vCenter.

Related concepts

Using object naming best practices when managing multiple clusters on page 49

Related tasks

Configuring vSphere Virtual Volumes (VVols) functionality on page 161

Viewing storage container details

You can review information for all active storage containers on the cluster in the NetApp Element Management extension point. You can also view details for each storage container, including efficiency and performance metrics and associated virtual volumes.

Before you begin

- You have enabled VVols functionality for the cluster.
- At least one storage container is available to select.

Steps

1. Select **NetApp Element Management > VVols**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Storage Containers** tab.
3. Select the check box for the storage container you want to review.
4. Click **Actions**.
5. In the resulting menu, select **Details**.

Related tasks

Configuring vSphere Virtual Volumes (VVols) functionality on page 161

Storage container details

You can view information about all active storage containers on the cluster on the **Storage Containers** page of the **VVols** tab from the **NetApp Element Management** extension point.

Account ID

The ID of the NetApp Element account associated with the storage container.

Name

The name of the storage container.

Status

The status of the storage container.

Possible values:

- `Active`: The storage container is in use.
- `Locked`: The storage container is locked.

Volume Count

The number of active volumes associated with the storage container account.

Individual storage container details

You can view storage container information for an individual storage container when you select it and view its details on the **Storage Containers** page of the **VVols** tab from the **NetApp Element Management** extension point.

The following details are in the **General Details** section:

Account ID

The ID of the cluster account associated with the storage container.

Storage container ID

The UUID of the virtual volume storage container.

Storage Container Name

The name of the associated storage container.

Datastore Name

The name of the associated datastore.

Status

The status of the storage container.

Possible values:

- `Active`: The storage container is in use.
- `Locked`: The storage container is locked.

Protocol Endpoint Type

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

Initiator Secret

The unique CHAP secret for the initiator.

Target Secret

The unique CHAP secret for the target.

Number of Volumes

The number of volumes associated with the storage container account.

The following details are in the **Efficiency** section:

Compression

The compression efficiency score for volumes in the account.

Deduplication

The deduplication efficiency score for volumes in the account.

Thin Provisioning

The thin provisioning efficiency score for volumes in the account.

Missing Volumes

The volumes that could not be queried for efficiency data.

Last Updated

The date and time of the last efficiency score.

The following details are in the **Performance Metrics** section:

Read Bytes

The total cumulative bytes read from all volumes in the account.

Read Operations

The total read operations to all volumes in the account since the creation of the account.

Write Bytes

The total cumulative bytes written to all volumes in the account.

Write Operations

The total cumulative write operations to all volume in the account since the creation of the account.

Unaligned Reads

For all 512e volumes in the account (virtual volumes are 512e by default), the number of read operations that were not on a 4k sector boundary. High numbers of unaligned reads may indicate improper partition alignment.

Unaligned Writes

For all 512e volumes in the account (virtual volumes are 512e by default), the number of write operations that were not on a 4k sector boundary. High numbers of unaligned writes may indicate improper partition alignment.

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.

Last Updated

The date and time of the last performance update.

The following details are in the **Virtual Volumes** section:

Volume ID

The ID of the underlying volume.

Virtual Volume ID

The UUID of the virtual volume.

Name

The name of the virtual volume.

Status

Status of the VVol task.

Deleting a storage container

You can delete storage containers from the NetApp Element Management extension point.

Before you begin

- An existing storage container is available to delete.
- All volumes have been removed from the storage container.

Steps

1. Select **NetApp Element Management > VVols**.

   Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Storage Containers** tab.
3. Select the check box for the storage container you want to delete.
4. Click **Actions**.
5. In the resulting menu, select **Delete**.
6. Confirm the action.
7. Refresh the list of storage containers in the **Storage Containers** sub-tab to confirm that the storage container has been removed.

Protocol endpoints

VMware ESXi hosts use logical I/O proxies known as protocol endpoints to communicate with virtual volumes. ESXi hosts bind virtual volumes to protocol endpoints to perform I/O operations. When a virtual machine on the host performs an I/O operation, the associated protocol endpoint directs I/O to the virtual volume with which it is paired.

Protocol endpoints in a NetApp Element cluster function as SCSI administrative logical units. Each protocol endpoint is created automatically by the cluster. For every node in a cluster, a corresponding protocol endpoint is created. For example, a four-node cluster will have four protocol endpoints.

iSCSI is the only supported protocol for NetApp Element software. Fibre Channel protocol is not supported.

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.

On the **VVols > Protocol Endpoints** page of the NetApp Element Management extension point, you can review protocol endpoint information.

Viewing protocol endpoint details

You can review information for all protocol endpoints on the cluster in the NetApp Element Management extension point.

Steps

1. Select **NetApp Element Management > VVols**.

Note: If two or more clusters are added, ensure that the cluster you intend to use for the task is selected in the navigation bar.

2. Click the **Protocol Endpoints** tab.

   The page displays information for all protocol endpoints on the cluster.

3. Select the check box for the protocol endpoint you want to review.
4. Click **Actions**.
5. In the resulting menu, select **Details**.

Protocol endpoint details

You can view information about all protocol endpoints on the cluster on the **Protocol Endpoint** page of the **VVols** tab from the **NetApp Element Management** extension point.

Primary Provider

The ID of the primary protocol endpoint provider.

Secondary Provider

The ID of the secondary protocol endpoint provider.

Protocol Endpoint ID

The UUID of the protocol endpoint.

Status

The status of the protocol endpoint.

Possible values:

- `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:

- `Primary`
- `Secondary`

SCSI NAA Device ID

The globally unique SCSI device identifier for the protocol endpoint in NAA IEEE Registered Extended Format.

Individual protocol endpoint details

You can view protocol endpoint information for an individual protocol endpoint when you select it and view its details on the **Protocol Endpoint** page of the **VVols** tab from the **NetApp Element Management** extension point.

Note: Virtual Volume details are only available when VMs are powered on.

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.

Status

The status of the protocol endpoint.

Possible values:

- `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:

- `Primary`
- `Secondary`

SCSI NAA Device ID

The globally unique SCSI device identifier for the protocol endpoint in NAA IEEE Registered Extended Format.

The following details are in the **Hosts** section:

Host Name

The name of the ESXi host.

Host Address

The IP address or DNS name for the ESXi host.

Initiator

Initiator IQNs for the virtual volume host.

Virtual Volume Host ID

The UUID for the ESXi host that hosts virtual volumes and is known to the cluster.

The following details are in the **Virtual Volumes** section:

Volume ID

The ID of the underlying volume.

Virtual Volume ID

The UUID of the virtual volume.

Name

The name of the virtual machine.

Status

Status of the VVol task.

Unregistering the vCenter Plug-in

You can unregister the NetApp Element Plug-in for vCenter Server using one of the procedures described for your installation. Unregistering the plug-in has the same effect as disabling it but does not remove all associated files and folders.

Before you begin

- vCenter Administrator role privileges to unregister a plug-in.
- IP address of the management node.
- URL and credentials for the vCenter from which you are unregistering the plug-in.

About this task

Note: Unregistering a plug-in package on vCenter Server does not delete the plug-in package files that are installed locally. To remove all plug-in files, see instructions on removing the plug-in.

Step

1. To unregister the plug-in, follow the procedure for your installed version:
   - `Version 3.0 or later`:
     - Unregister the plug-in using the vCenter Plug-in registration utility:
       a. Enter the IP address for your management node in a browser, including the TCP port for registration: `https://[management node IP]:9443`.
       b. Click or navigate to **Unregister Plug-in**.
       c. Enter the following:
          i. The IP address or FQDN server name of the vCenter service on which you have registered your plug-in.
          ii. The vCenter Administrator user name.
          iii. The vCenter Administrator password.
       d. Click **Unregister**.
   - `For versions earlier than 3.0 to 2.7 :`
     - Use the vCenter Managed Object Browser (MOB) interface in your browser to manually unregister.
       a. Enter the MOB URL: `https://[vcenter]/mob`
       b. Click **Content** > **Extension Manager** > **UnregisterExtension**.
       c. Enter `com.solidfire`.
       d. Click **Invoke Method**.
     - Unregister using PowerCLI:

```
Connect-VIServer -Server $vcenter -User
administrator@vsphere.local
-Password xxxxXXx -Force -ErrorAction Stop -SaveCredentials
```

```
$em = Get-View ExtensionManager
$em.ExtensionList | ft -property Key
$em.UnregisterExtension("com.solidfire")
$em.UpdateViewData()
$em.ExtensionList | ft -property Key
Disconnect-VIServer * -Confirm:$false
```

- `Version 2.6.1 or earlier`:

  a. In a browser, enter the URL for the registration utility or locate it in the program directory:

  - `https://<FDVA or management node IP>:8443`
  - `/opt/solidfire/vcp/bin/vcp-reg.sh`

  b. In the vCenter Plugin Register/Unregister window, click **Unregister**.

Related tasks

Removing the vCenter Plug-in on page 177

Removing the vCenter Plug-in

For vCenter Plug-in 2.5 or later, you must manually remove files from vCenter Server.

Before you begin

- You have unregistered the existing plug-in and have SSH, RDP, or other appropriate connectivity to vCSA or vCenter Server.

About this task

You must complete the following process to remove all files associated with the plug-in.

Steps

1. Log in as an administrator to the server that is running vCenter Server and open a command prompt.
2. Stop vCenter Server services:
   - Windows:
     - (For Flash clients) Enter the following command:

       ```
       C:\Program Files\VMware\vCenter Server\vmon>.\vmon-cli --stop
       vsphere-client
       ```

     - (For HTML5 clients) Enter the following commands:

       ```
       C:\Program Files\VMware\vCenter Server\vmon>.\vmon-cli --stop
       vsphere-client
       C:\Program Files\VMware\vCenter Server\vmon>.\vmon-cli --stop
       vsphere-ui
       ```

   - vCenter Server Appliance (vCSA): Use the following commands:
     - (For Flash clients) Enter the following command:

       ```
       service-control --stop vsphere-client
       ```

     - (For HTML5 clients) Enter the following commands:

       ```
       service-control --stop vsphere-client
       service-control --stop vsphere-ui
       ```

3. Remove SolidFire folders and files from the following locations:
   - Windows: Use Windows Explorer and search for “SolidFire” in `C:\ProgramData\VMware` and `C:\Program Files\VMware`.

     Note: The ProgramData folder is hidden. You must enter the complete file path to access the folder.

   - vCSA: Use the following command:

     ```
     find / -name "*solidfire*" -exec rm -rf {} \;
     ```

4. Start vCenter Server services:
   - Windows:
     - (For Flash clients) Enter the following command:

       ```
       C:\Program Files\VMware\vCenter Server\vmon>.\vmon-cli --start
       vsphere-client
       ```

     - (For HTML5 clients) Enter the following commands:

       ```
       C:\Program Files\VMware\vCenter Server\vmon>.\vmon-cli --start
       vsphere-client
       C:\Program Files\VMware\vCenter Server\vmon>.\vmon-cli --start
       vsphere-ui
       ```

   - vCSA: Use the following commands:
     - (For Flash clients) Enter the following command:

       ```
       service control --start vsphere-client
       ```

     - (For HTML5 clients) Enter the following commands:

       ```
       service control --start vsphere-client
       service-control --start vsphere-ui
       ```

Related tasks

Unregistering the vCenter Plug-in on page 175

Troubleshooting

You need to be aware of some of the common issues with the NetApp Element Plug-in for vCenter Server and the steps to resolve them.

Plug-in registration successful but icons do not appear in web client

Description

Registration shows as successful, but the plug-in icons are not visible from the vSphere Web Client.

Corrective action

- Log out of the vSphere Web Client and log in again. Closing and re-opening your browser may be required.
- Clear your browser cache.
- From vCenter, restart the vSphere Web Client Service from the Services menu within Windows Administrative Tools or reboot vCenter.
- Ensure that you have all required default administrative privileges associated with the vCenter Administrator role.
- Check that the plug-in ZIP file successfully downloaded to vCenter:

  1. Open `vsphere_client_virgo.log` in the vCenter. vCenter log files for versions 6.5 and 6.7 are in the following locations:
     - Flash installations: `/var/log/vmware/vsphere-client/logs/vsphere_client_virgo.log`
     - HTML5 installations: `/var/log/vmware/vsphere-ui/logs/vsphere_client_virgo.log`
  2. If a failure message indicates that the ZIP download failed, download the ZIP again.

     Note: You might need to correct an unreachable or bad URL. Update the plug-in registration or unregister and register the plug-in again with a corrected URL. Failure to download the ZIP can also occur if you specified an HTTP URL without changing the `allowHTTP` setting.

- Verify networking ports. Ensure the management node is reachable from vCenter bidirectionally on the required ports.
- Check the vCenter's MOB extension record (`https://<vcenterIP>/mob/?moid=ExtensionManager&doPath=extensionList%5b%22com%2esolidfire%22%5d%2eserver`) that contains the download location URL for the plug-in ZIP:

  1. Paste the URL into a browser.
  2. Verify that the plug-in ZIP can be downloaded.
     - If the plug-in ZIP can be downloaded, proceed to the next step.

- If the plug-in ZIP cannot be downloaded, check for networking issues between vCenter Server and the management node.

3. If the plug-in cannot be downloaded, compare the `serverThumbprint` in the MOB record with the certificate SHA-1 for the ZIP URL that is displayed in the browser:
   a. If the registration record in the MOB has an incorrect or stale URL or SHA-1, unregister the plug-in and register the plug-in again.
   b. If the problem persists and the ZIP is unreachable, inspect the ZIP URL to determine if there is an issue with the management node address used. In some cases, it might be necessary to customize a URL using the registration utility for the plug-in so that the ZIP file can be downloaded.

Related tasks

Modifying vCenter properties for an in-house (dark site) HTTP server on page 25

Related references

Network port requirements on page 10

Error registering plug-in using Registration UI

Description

When using the registration utility, there is an error registering the plug-in against the vCenter server. A plug-in with the key `com.solidfire` is already installed.

Corrective action

In the registration utility, use **Update Plug-in** instead of **Register Plug-in**.

Error updating plug-in using Registration UI

Description

When using the registration utility, there is an error updating the plug-in against the vCenter server. A plug-in with the key `com.solidfire` is not installed for the update.

Corrective action

In the registration utility, use **Register Plug-in** instead of **Update Plug-in**.

Keystore error using Registration UI

Description

When using the registration utility, there is an error that indicates the keystore cannot be found at `/opt/solidfire/registration/keystore` and `/var/cache/jetty/tmp/./etc/keystore`.

Corrective action

1. Reboot the management node or execute the following:

```
sudo /opt/solidfire/vcp-reg.bash -F
```

2. Refresh the registration utility web UI.

Error message that NetApp extension cannot be upgraded

Message

```
org.springframework.transaction.CannotCreateTransactionException:
Could not open JPA EntityManager for transaction; nested exception
is javax.persistence.PersistenceException:
org.hibernate.exception.GenericJDBCException: Could not open
connection.
```

Description

During a Windows vCenter Server upgrade from version 6.0 to 6.5, you see a warning that the NetApp Extension cannot be upgraded or may not work with the new vCenter Server. After you complete the upgrade and log in to the vSphere Web Client, the error occurs when you select a vCenter Plug-in extension point. This error occurs because the directory that stores the runtime database has changed from version 6.0 to 6.5. The vCenter Plug-in is unable to create the needed files for runtime.

Corrective action

1. Unregister the plug-in.
2. Remove plug-in files.
3. Reboot the vCenter.
4. Register the plug-in.
5. Log in to the vSphere Web Client.

Related tasks

Unregistering the vCenter Plug-in on page 175
Removing the vCenter Plug-in on page 177
Registering the vCenter Plug-in with vCenter on page 22

Removing plug-in completes successfully but icons remain

Description

Removing vCenter Plug-in package files completed successfully, but plug-in icons are still visible in the vSphere Web Client.

Corrective action

Log out of the vSphere Web Client and log in again. Closing and re-opening your browser might be required. If logging out of vSphere Web Client does not resolve the issue, it might be necessary to reboot the vCenter server web services.

Plug-in cannot be unregistered or removed after admin password change

Description

After the admin password for the vCenter that was used to register the plug-in is changed, the vCenter Plug-in cannot be unregistered or removed.

Corrective action

For plug-in 2.6, go to the vCenter Plug-in Register/Unregister page. Click the **Update** button to change the vCenter IP address, user ID, and password.

For plug-in 2.7 or later, update the vCenter Administrator password in **mNode Settings** in the plug-in.

Plug-in management tasks fail or volumes are not accessible to ESXi host

Description

Create, clone, and share datastore tasks fail or volumes are not accessible by the ESXi host.

Corrective action

- Check that the software iSCSI HBA is present and enabled on the ESXi host for datastore operations.
- Check that the volume is not deleted or assigned to an incorrect volume access group.
- Check that the volume access group has the correct host IQN.
- Check that the associated account has the correct CHAP settings.
- Check that volume status is `active`, volume access is `readWrite`, and 512e is set to `true`.

Failure occurs during vCenter Plug-in use on Firefox 59.0.2 browsers

Message

```
Name:HttpErrorResponse Raw Message:Http failure response for
https://vc6/ui/solidfire-war-4.2.0-SNAPSHOT/rest/vsphere//servers:
500 Internal Server Error Return Message:Server error. Please try
again or contact NetApp support
```

Description

This issue occurs in vSphere HTML5 web clients using Firefox. The vSphere Flash client is not affected.

Corrective action

Use the full FQDN in the browser URL.

Delete datastore operation fails

Description

A delete datastore operation fails.

Corrective action

Check that all VMs have been deleted from the datastore. You must delete VMs from a datastore before the datastore can be deleted.

Cluster pair cannot connect using a pairing key

Description

A connection error occurs during cluster pairing using a pairing key. The error message in the **Create Cluster Pairing** dialog box indicates that there is no route to host.

Corrective action

Manually delete the unconfigured cluster pair the process created on the local cluster and perform the cluster pairing again.

Error message for QoSSIOC status

Description

QoSSIOC status for the plug-in displays a warning icon and error message.

Corrective action

- `Unable to reach IP address`

  The IP address is invalid or no responses are received. Verify that the address is correct and that the management node is online and available.

- `Unable to communicate`

  The IP address can be reached but calls to the address fail. This might indicate that the QoSSIOC service is not running at the specified address or a firewall might be blocking traffic.

- `Unable to connect to the SIOC service`

  Open `sioc.log` in `/var/log` on the management node (`/var/log/solidfire/` on older management nodes) to verify that the SIOC service started successfully. SIOC service startup can take 50 seconds or more. If the service did not start successfully, try again. You can verify the current status of the SIOC service by opening `siocStatus.log` in `/var/log/` on the management node.

Related tasks

Configuring management node settings for QoSSIOC on page 58

Related references

Network port requirements on page 10

QoSSIOC service shown as available but is unavailable

Description

QoSSIOC service settings displays as **UP**, but QoSSIOC is unavailable.

Corrective action

From the **mNode Settings** tab in the NetApp Element Configuration extension point, click the refresh button in the QoSSIOC Service. Update the IP address or user authentication information as needed.

QoSSIOC is enabled for datastore but unavailable

Description

QoSSIOC is enabled for a datastore, but QoSSIOC is unavailable.

Corrective action

Check that the VMware SIOC is enabled on the datastore:

1. Open `sioc.log` in `/var/log` on the management node.
2. Search for this text:

   ```
   SIOC is not enabled
   ```

3. Use the vSphere Web Client or CLI to enable SIOC on the datastore.

Where to find product documentation and other information

You can learn more about using and managing NetApp HCI and SolidFire all-flash storage from the resources available in the Documentation Centers and Resources pages for both products.

In the Documentation Centers, you can also find information about hardware installation and maintenance, additional content resources available, links to known issues and resolved issues, and the latest release notes. On the Resources pages, you can find links to data sheets, technical reports, white papers, and videos.

- *NetApp HCI Documentation Center*
- *NetApp HCI Resources page*
- *SolidFire and Element 11.5 Documentation Center*
- *SolidFire and Element 11.3 Documentation Center*
- *SolidFire Resources page*

Contacting NetApp Support

If you need help with or have questions or comments about NetApp products, contact NetApp Support.

- Web:
  mysupport.netapp.com
- Phone:
  - 888.4.NETAPP (888.463.8277) (US and Canada)
  - 00.800.44.638277 (EMEA/Europe)
  - +800.800.80.800 (Asia/Pacific)

Copyright

Copyright © 2019 NetApp, Inc. All rights reserved. Printed in the U.S.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark 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

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., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277