



Configure cluster settings

Element Software

NetApp

February 03, 2022

Table of Contents

- Configure cluster settings 1
 - Find more information 1
 - Enable and disable encryption at rest for a cluster 1
 - Set the cluster full threshold 3
 - Enable and disable support access 3
 - Manage the Terms of Use banner 3
 - Set the Network Time Protocol 4
 - Manage SNMP 6
 - Manage drives 8
 - Manage nodes 9
 - View Fibre Channel ports details 12
 - Manage virtual networks 13

Configure cluster settings

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

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

Options

- [Work with virtual volumes](#)
- [Use SnapMirror replication between Element and ONTAP clusters](#)
- [Set the cluster full threshold](#)
- [Enable and disable support access](#)
- [How are the blockSpace thresholds calculated for Element](#)
- [Enable and disable encryption for a cluster](#)
- [Manage the Terms of Use banner](#)
- [Configure Network Time Protocol servers for the cluster to query](#)
- [Manage SNMP](#)
- [Manage drives](#)
- [Manage nodes](#)
- [Manage virtual networks](#)
- [View Fibre Channel ports details](#)

Find more information

- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

Enable and disable encryption at rest for a cluster

With SolidFire clusters, you can encrypt all at-rest data stored on cluster drives. You can enable cluster-wide protection of self-encrypting drives (SED) using either [hardware or software-based encryption at rest](#).

You can enable hardware encryption at rest using the Element UI or API. Enabling the hardware encryption at rest feature does not affect performance or efficiency on the cluster. You can enable software encryption at rest using the Element API only.

Hardware-based encryption at rest is not enabled by default during cluster creation and can be enabled and disabled from the Element UI.



For SolidFire all-flash storage clusters, software encryption at rest must be enabled during cluster creation and cannot be disabled after the cluster has been created. For SolidFire Enterprise SDS (eSDS) clusters, software encryption at rest is enabled by default.

What you'll need

- You have cluster administrator privileges to enable or change encryption settings.
- For hardware-based encryption at rest, you have ensured that the cluster is in a healthy state before changing encryption settings.
- If you are disabling encryption, two nodes must be participating in a cluster to access the key to disable encryption on a drive.

Check encryption at rest status

To see the current status of encryption at rest and/or software encryption at rest on the cluster, use the [GetClusterInfo](#) method. You can use the [GetSoftwareEncryptionAtRestInfo](#) method to get information the cluster uses to encrypt data at rest.



The Element software UI dashboard at <https://<MVIP>/> currently only shows encryption at rest status for hardware-based encryption.

Options

- [Enable hardware-based encryption at rest](#)
- [Enable software-based encryption at rest](#)
- [Disable hardware-based encryption at rest](#)

Enable hardware-based encryption at rest



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

1. From the Element UI, select **Cluster > Settings**.
2. Select **Enable Encryption at Rest**.

Enable software-based encryption at rest



Software encryption at rest cannot be disabled after it is enabled on the cluster.

1. During cluster creation, run the [create cluster method](#) with `enableSoftwareEncryptionAtRest` set to `true`.

Disable hardware-based encryption at rest

1. From the Element UI, select **Cluster > Settings**.
2. Select **Disable Encryption at Rest**.

Find more information

- [NetApp SolidFire Resources Page](#)
- [Documentation for earlier versions of NetApp SolidFire and Element products](#)

Set the cluster full threshold

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

What you'll need

You must have cluster administrator privileges.

Steps

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

Find more information

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

Enable and disable support access

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

You must have cluster admin privileges to change support access.



This feature is unavailable in SolidFire Enterprise SDS clusters.

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

Manage the Terms of Use banner

You can enable, edit, or configure a banner that contains a message for the user.

Options

[Enable the Terms of Use banner](#) [Edit the Terms of Use banner](#) [Disable the Terms of Use banner](#)

Enable the Terms of Use banner

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

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

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



Do not exceed 4096 characters.

3. Click **Enable**.

Edit the Terms of Use banner

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

What you'll need

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

Steps

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



Do not exceed 4096 characters.

3. Click **Save Changes**.

Disable the Terms of Use banner

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

What you'll need

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

Steps

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

Set the Network Time Protocol

Setting up the Network Time Protocol (NTP) can be achieved in one of two ways: either instruct each node in a cluster to listen for broadcasts or or instruct each node to query an NTP server for updates.

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



This feature is unavailable in SolidFire Enterprise SDS clusters.

Configure Network Time Protocol servers for the cluster to query

You can instruct each node in a cluster to query a Network Time Protocol (NTP) server for updates. The cluster contacts only configured servers and requests NTP information from them.

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

Using the FQDN depends on whether the individual storage node's DNS settings are in place and operational. To do so, configure the DNS servers on every storage node and ensure that the ports are open by reviewing the [Network Port Requirements](#) page.

You can enter up to five different NTP servers.



You can use both IPv4 and IPv6 addresses.

What you'll need

You must have cluster administrator privileges to configure this setting.

Steps

1. Configure a list of IPs and/or FQDNs in the server settings.
2. Ensure that the DNS is set properly on the nodes.
3. Click **Cluster > Settings**.
4. Under Network Time Protocol Settings, select **No**, which uses the standard NTP configuration.
5. Click **Save Changes**.

Find more information

- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

Configure the cluster to listen for NTP broadcasts

By using the broadcast mode, you can instruct each node in a cluster to listen on the network for Network Time Protocol (NTP) broadcast messages from a particular server.

What you'll need

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

Steps

1. Click **Cluster > Settings**.
2. Enter the NTP server or servers that are using broadcast mode into the server list.
3. Under Network Time Protocol Settings, select **Yes** to use a broadcast client.
4. To set the broadcast client, in the **Server** field, enter the NTP server you configured in broadcast mode.

5. Click **Save Changes**.

Find more information

- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

Manage SNMP

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

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



You can use both IPv4 and IPv6 addresses.

SNMP details

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

- **SNMP MIBs**

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

- **General SNMP Settings**

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

- **SNMP Trap Settings**

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

Configure an SNMP requestor

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

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



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

5. **Optional:** To add another requestor, click **Add a Requestor** and enter the **Community String** and

Network information.

6. Click **Save Changes**.

Find more information

- [Configure SNMP traps](#)
- [View managed object data using management information base files](#)

Configure an SNMP USM user

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

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

Configure SNMP traps

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

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

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

View managed object data using management information base files

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

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

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

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

Steps

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

Manage drives

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

For more information

- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

Drives details

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

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

The following elements appear in the list of active drives.

- **Drive ID**

The sequential number assigned to the drive.

- **Node ID**

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

- **Node Name**

The name of the node that houses the drive.

- **Slot**

The slot number where the drive is physically located.

- **Capacity**

The size of the drive, in GB.

- **Serial**

The serial number of the drive.

- **Wear Remaining**

The wear level indicator.

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

- **Type**

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

Manage nodes

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

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

Find more information

[Add a node to a cluster](#)

Add a node to a cluster

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

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

You can add nodes of smaller or larger capacities to an existing cluster. You can add larger node capacities to

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



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

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

You can add nodes to SolidFire Enterprise SDS (SolidFire eSDS) clusters or to NetApp HCI appliances.

SolidFire eSDS nodes are denoted with a Node type prefix of "SFc", for example, "SFc100."

Steps

1. Select **Cluster > Nodes**.
2. Click **Pending** to view the list of pending nodes.

When the process for adding nodes (both SolidFire eSDS or non SolidFire eSDS nodes) is complete, they appear in the Active nodes list. Until then, pending nodes appear in the Pending Active list.

For nodes that are not SolidFire eSDS nodes, SolidFire installs the Element software version of the cluster on the pending nodes when you add them to a cluster. This might take a few minutes.

3. Do one of the following:
 - To add individual nodes, click the **Actions** icon for the node you want to add.
 - To add multiple nodes, select the check box of the nodes to add, and then **Bulk Actions**. **Note:** If the node you are adding has a different version of Element software than the version running on the cluster, the cluster asynchronously updates the node to the version of Element software running on the cluster master. After the node is updated, it automatically adds itself to the cluster. During this asynchronous process, the node will be in a pendingActive state.
4. Click **Add**.

The node appears in the list of active nodes.

Find more information

[Node versioning and compatibility](#)

Node versioning and compatibility

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

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

- **Major**

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

- **Minor**

The second number designates smaller software features or enhancements to existing software features that have been added to a major release. This component is incremented within a major version component to indicate that this incremental release is not compatible with any other Element software incremental releases with a different minor component. For example, 11.0 is not compatible with 11.1, and 11.1 is not compatible with 11.2.

- **Micro**

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

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

Cluster capacity in a mixed node environment

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

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

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

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

- SF3010
- SF6010
- SF9010

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

View node details

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

Steps

1. Click **Cluster > Nodes**.
2. To view the details for a specific node, click the **Actions** icon for a node.
3. Click **View Details**.
4. Review the node details:
 - **Node ID**: The system-generated ID for the node.
 - **Node Name**: The host name for the node.
 - **Available 4k IOPS**: The IOPS configured for the node.
 - **Node Role**: The role that the node has in the cluster. Possible values:
 - **Cluster Master**: The node that performs cluster-wide administrative tasks and contains the MVIP and SVIP.
 - **Ensemble Node**: A node that participates in the cluster. There are either 3 or 5 ensemble nodes depending on cluster size.
 - **Fibre Channel**: A node in the cluster.
 - **Node Type**: The model type of the node.
 - **Active Drives**: The number of active drives in the node.
 - **Management IP**: The management IP (MIP) address assigned to node for 1GbE or 10GbE network admin tasks.
 - **Cluster IP**: The cluster IP (CIP) address assigned to the node used for the communication between nodes in the same cluster.
 - **Storage IP**: The storage IP (SIP) address assigned to the node used for iSCSI network discovery and all data network traffic.
 - **Management VLAN ID**: The virtual ID for the management local area network.
 - **Storage VLAN ID**: The virtual ID for the storage local area network.
 - **Version**: The version of software running on each node.
 - **Replication Port**: The port used on nodes for remote replication.
 - **Service Tag**: The unique service tag number assigned to the node.

View Fibre Channel ports details

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

View information about the Fibre Channel ports that are connected to the cluster.

Steps

1. Click **Cluster > FC Ports**.
2. To filter information on this page, click **Filter**.
3. Review the details:
 - **Node ID**: The node hosting the session for the connection.
 - **Node Name**: System-generated node name.
 - **Slot**: Slot number where the Fibre Channel port is located.

- **HBA Port:** Physical port on the Fibre Channel host bus adapter (HBA).
- **WWNN:** The world wide node name.
- **WWPN:** The target world wide port name.
- **Switch WWN:** World wide name of the Fibre Channel switch.
- **Port State:** Current state of the port.
- **nPort ID:** The node port ID on the Fibre Channel fabric.
- **Speed:** The negotiated Fibre Channel speed. Possible values are as follows:
 - 4Gbps
 - 8Gbps
 - 16Gbps

Find more information

- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

Manage virtual networks

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

Find more information

- [Add a virtual network](#)
- [Enable virtual routing and forwarding](#)
- [Edit a virtual network](#)
- [Edit VRF VLANs](#)
- [Delete a virtual network](#)

Add a virtual network

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

What you'll need

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



You must consider the following criteria for this configuration:

- VLANs that are not VRF-enabled require initiators to be in the same subnet as the SVIP.

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

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. Click **Cluster > Network**.
2. Click **Create VLAN**.
3. In the **Create a New VLAN** dialog box, enter values in the following fields:
 - **VLAN Name**
 - **VLAN Tag**
 - **SVIP**
 - **Netmask**
 - (Optional) **Description**
4. Enter the **Starting IP** address for the range of IP addresses in **IP Address Blocks**.
5. Enter the **Size** of the IP range as the number of IP addresses to include in the block.
6. Click **Add a Block** to add a non-continuous block of IP addresses for this VLAN.
7. Click **Create VLAN**.

View virtual network details

Steps

1. Click **Cluster > Network**.
2. Review the details.
 - **ID**: Unique ID of the VLAN network, which is assigned by the system.
 - **Name**: Unique user-assigned name for the VLAN network.
 - **VLAN Tag**: VLAN tag assigned when the virtual network was created.
 - **SVIP**: Storage virtual IP address assigned to the virtual network.
 - **Netmask**: Netmask for this virtual network.
 - **Gateway**: Unique IP address of a virtual network gateway. VRF must be enabled.
 - **VRF Enabled**: Indication of whether virtual routing and forwarding is enabled or not.
 - **IPs Used**: The range of virtual network IP addresses used for the virtual network.

Enable virtual routing and forwarding

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

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

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

Find more information

[Add a virtual network](#)

Edit a virtual network

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

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

When managing the size of VLAN IP address ranges, you should note the following limitations:

- You can only remove IP addresses from the initial IP address range assigned at the time the VLAN was created.
- You can remove an IP address block that was added after the initial IP address range, but you cannot resize an IP block by removing IP addresses.
- When you try to remove IP addresses, from either the initial IP address range or in an IP block, that are in use by nodes in the cluster, the operation might fail.
- You cannot reassign specific in-use IP addresses to other nodes in the cluster.

You can add an IP address block by using the following procedure:

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

Link to troubleshooting KB articles

Link to the Knowledge Base articles for help with troubleshooting issues with managing your VLAN IP address ranges.

- [Duplicate IP warning after adding a storage node in VLAN on Element cluster](#)
- [How to determine which VLAN IP's are in use and which nodes those IP's are assigned to in Element](#)

Edit VRF VLANs

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

1. Click **Cluster > Network**.
2. Click the Actions icon for the VLAN you want to edit.
3. Click **Edit**.
4. Enter the new attributes for the VRF VLAN in the **Edit VLAN** dialog box.
5. Click **Save Changes**.

Delete a virtual network

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

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

Find more information

[Edit a virtual network](#)

Copyright Information

Copyright © 2022 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means-graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system- without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

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

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.