



# Manage your system

## Element Software

NetApp  
February 03, 2022

# Table of Contents

- Manage your system ..... 1
  - For more information ..... 1
  - Enable multi-factor authentication ..... 1
  - Configure cluster settings ..... 2
  - Create a cluster supporting FIPS drives ..... 18
  - Enable FIPS 140-2 for HTTPS on your cluster ..... 21
  - Get started with external key management ..... 23

# Manage your system

You can manage your system in the Element UI. This includes enabling multi-factor authentication, managing cluster settings, supporting Federal Information Processing Standards (FIPS), and using external key management.

- [Enable multi-factor authentication](#)
- [Configure cluster settings](#)
- [Create a cluster supporting FIPS drives](#)
- [Get started with external key management](#)

## For more information

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

## Enable multi-factor authentication

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

### Set up multi-factor authentication

You can use these basic steps via the Element API to set up your cluster to use multi-factor authentication.

Details of each API method can be found in the [Element API Reference](#).

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

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

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

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

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



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

- email=[bob@company.com](#) — where the IdP is configured to release an email address in the SAML attributes.
- group=cluster-administrator - where the IdP is configured to release a group property in which all users should have access. Note that the SAML attribute Name/Value pairing is case-sensitive for security purposes.

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

### Find more information

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

## Additional information for multi-factor authentication

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

- In order to refresh IdP certificates that are no longer valid, you will need to use a non-IdP admin user to call the following API method: `UpdateIdpConfiguration`
- MFA is incompatible with certificates that are less than 2048 bits in length. By default, a 2048-bit SSL certificate is created on the cluster. You should avoid setting a smaller sized certificate when calling the API method: `SetSSLCertificate`



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

- IdP admin users cannot be used to make API calls directly (for example, via SDKs or Postman) or used for other integrations (for example, OpenStack Cinder or vCenter Plug-in). Add either LDAP cluster admin users or local cluster admin users if you need to create users that have these abilities.

### Find more information

- [Managing storage with the Element API](#)
- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

## 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](#)

## Create a cluster supporting FIPS drives

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

- [Avoid mixing nodes for FIPS drives](#)
- [Enable encryption at rest](#)
- [Identify whether nodes are ready for the FIPS drives feature](#)
- [Enable the FIPS drives feature](#)
- [Check the FIPS drive status](#)
- [Troubleshoot the FIPS drive feature](#)

### Avoid mixing nodes for FIPS drives

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

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

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

### Enable encryption at rest

You can enable and disable cluster-wide encryption at rest. This feature is not enabled by default. To support FIPS drives, you must enable encryption at rest.

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

**Find more information**

- [Enable and disable encryption for a cluster](#)
- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

## Identify whether nodes are ready for the FIPS drives feature

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

The resulting report shows one of the following statuses:

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

### Steps

1. Using the Element API, check to see if the nodes and drives in the storage cluster are capable of FIPS drives by entering:

```
GetFipsReport
```

2. Review the results, noting any nodes that did not display a status of Ready.
3. For any nodes that did not display a Ready status, check to see if the drive is capable of supporting the FIPS drives feature:
  - Using the Element API, enter: `GetHardwareList`
  - Note the value of the **DriveEncryptionCapabilityType**. If it is "fips," the hardware can support the FIPS drives feature.

See details about `GetFipsReport` or `ListDriveHardware` in the [Element API Reference](#).

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

### Find more information

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

## Enable the FIPS drives feature

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

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

### Step

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

```
EnableFeature params: FipsDrives
```

## Find more information

- [Manage storage with the Element API](#)
- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

## Check the FIPS drive status

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

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

```
GetFeatureStatus
```

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

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

## Find more information

- [Manage storage with the Element API](#)
- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

## Troubleshoot the FIPS drive feature

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

1. Using the Element UI, select **Reporting > Alerts**.
2. Look for cluster faults including:
  - FIPS drives mismatched
  - FIPS drives out of compliance
3. For resolution suggestions, see Cluster Fault code information.

## Find more information

- [Cluster fault codes](#)
- [Manage storage with the Element API](#)
- [SolidFire and Element Resources page](#)

## Enable FIPS 140-2 for HTTPS on your cluster

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

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



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

For more information, see the Element API information.

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

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

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

### Find more information

- [SSL ciphers](#)
- [Manage storage with the Element API](#)
- [SolidFire and Element Resources page](#)
- [NetApp Element Plug-in for vCenter Server](#)

### SSL ciphers

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

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

- **FIPS 140-2 disabled**

TLS\_DHE\_RSA\_WITH\_AES\_128\_CBC\_SHA256 (dh 2048) - A  
TLS\_DHE\_RSA\_WITH\_AES\_128\_GCM\_SHA256 (dh 2048) - A  
TLS\_DHE\_RSA\_WITH\_AES\_256\_CBC\_SHA256 (dh 2048) - A  
TLS\_DHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (dh 2048) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_128\_CBC\_SHA256 (secp256r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_128\_GCM\_SHA256 (secp256r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_256\_CBC\_SHA384 (secp256r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (secp256r1) - A  
TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA (rsa 2048) - C  
TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA256 (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_128\_GCM\_SHA256 (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA256 (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (rsa 2048) - A  
TLS\_RSA\_WITH\_CAMELLIA\_128\_CBC\_SHA (rsa 2048) - A  
TLS\_RSA\_WITH\_CAMELLIA\_256\_CBC\_SHA (rsa 2048) - A  
TLS\_RSA\_WITH\_IDEA\_CBC\_SHA (rsa 2048) - A  
TLS\_RSA\_WITH\_RC4\_128\_MD5 (rsa 2048) - C  
TLS\_RSA\_WITH\_RC4\_128\_SHA (rsa 2048) - C  
TLS\_RSA\_WITH\_SEED\_CBC\_SHA (rsa 2048) - A

- **FIPS 140-2 enabled**

TLS\_DHE\_RSA\_WITH\_AES\_128\_CBC\_SHA256 (dh 2048) - A  
TLS\_DHE\_RSA\_WITH\_AES\_128\_GCM\_SHA256 (dh 2048) - A  
TLS\_DHE\_RSA\_WITH\_AES\_256\_CBC\_SHA256 (dh 2048) - A  
TLS\_DHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (dh 2048) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_128\_CBC\_SHA256 (sect571r1) - A



TLS\_ECDHE\_RSA\_WITH\_AES\_128\_CBC\_SHA256 (secp256r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_128\_GCM\_SHA256 (secp256r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_128\_GCM\_SHA256 (sect571r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_256\_CBC\_SHA384 (sect571r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_256\_CBC\_SHA384 (secp256r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (secp256r1) - A  
TLS\_ECDHE\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (sect571r1) - A  
TLS\_RSA\_WITH\_3DES\_EDE\_CBC\_SHA (rsa 2048) - C  
TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_128\_CBC\_SHA256 (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_128\_GCM\_SHA256 (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_256\_CBC\_SHA256 (rsa 2048) - A  
TLS\_RSA\_WITH\_AES\_256\_GCM\_SHA384 (rsa 2048) - A

### Find more information

[Enable FIPS 140-2 for HTTPS on your cluster](#)

## Get started with external key management

External key management (EKM) provides secure Authentication Key (AK) management in conjunction with an off-cluster external key server (EKS). The AKs are used to lock and unlock Self Encrypting Drives (SEDs) when [encryption at rest](#) is enabled on the cluster. The EKS provides secure generation and storage of the AKs. The cluster utilizes the Key Management Interoperability Protocol (KMIP), an OASIS defined standard protocol, to communicate with the EKS.



Only software encryption at rest is available for SolidFire Enterprise SDS clusters.

- [Set up external management](#)
- [Rekey software encryption at rest master key](#)
- [Recover inaccessible or invalid authentication keys](#)
- [External key management API commands](#)

### Find more information

- [CreateCluster API that can be used to enable software encryption at rest](#)

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

## Set up external key management

You can follow these steps and use the Element API methods listed to set up your external key management feature.

### What you'll need

- If you are setting up external key management in combination with software encryption at rest, you have enabled software encryption at rest using the [CreateCluster](#) method on a new cluster that does not contain volumes.

### Steps

1. Establish a trust relationship with the External Key Server (EKS).
  - a. Create a public/private key pair for the Element cluster that is used to establish a trust relationship with the key server by calling the following API method: [CreatePublicPrivateKeyPair](#)
  - b. Get the certificate sign request (CSR) which the Certification Authority needs to sign. The CSR enables the key server to verify that the Element cluster that will be accessing the keys is authenticated as the Element cluster. Call the following API method: [GetClientCertificateSignRequest](#)
  - c. Use the EKS/Certificate Authority to sign the retrieved CSR. See third-party documentation for more information.
2. Create a server and provider on the cluster to communicate with the EKS. A key provider defines where a key should be obtained, and a server defines the specific attributes of the EKS that will be communicated with.
  - a. Create a key provider where the key server details will reside by calling the following API method: [CreateKeyProviderKmip](#)
  - b. Create a key server providing the signed certificate and the public key certificate of the Certification Authority by calling the following API methods: [CreateKeyServerKmip](#) [TestKeyServerKmip](#)

If the test fails, verify your server connectivity and configuration. Then repeat the test.
  - c. Add the key server into the key provider container by calling the following API methods: [AddKeyServerToProviderKmip](#) [TestKeyProviderKmip](#)

If the test fails, verify your server connectivity and configuration. Then repeat the test.
3. Do one of the following as a next step for encryption at rest:
  - a. (For hardware encryption at rest) Enable [hardware encryption at rest](#) by providing the ID of the key provider that contains the key server used for storing the keys by calling the [EnableEncryptionAtRest](#) API method.



You must enable encryption at rest via the [API](#). Enabling encryption at rest using the existing Element UI button will cause the feature to revert to using internally generated keys.

- b. (For software encryption at rest) In order for [software encryption at rest](#) to utilize the newly created key provider, pass the key provider ID to the [RekeySoftwareEncryptionAtRestMasterKey](#) API method.

## Find more information

- [Enable and disable encryption for a cluster](#)
- [NetApp SolidFire Resources Page](#)
- [Documentation for earlier versions of NetApp SolidFire and Element products](#)

## Rekey software encryption at rest master key

You can use the Element API to rekey an existing key. This process creates a new replacement master key for your external key management server. Master keys are always replaced by new master keys and never duplicated or overwritten.

You might need to rekey as part of one of the following procedures:

- Create a new key as part of a change from internal key management to external key management.
- Create a new key as a reaction to or as protection against a security-related event.



This process is asynchronous and returns a response before the rekey operation is complete. You can use the [GetAsyncResult](#) method to poll the system to see when the process has completed.

## What you'll need

- You have enabled software encryption at rest using the [CreateCluster](#) method on a new cluster that does not contain volumes and has no I/O. Use [GetSoftwareEncryptionAtRestInfo](#) to confirm that the state is enabled before proceeding.
- You have [established a trust relationship](#) between the SolidFire cluster and an External Key Server (EKS). Run the [TestKeyProviderKmpip](#) method to verify that a connection to the key provider is established.

## Steps

1. Run the [ListKeyProvidersKmpip](#) command and copy the key provider ID (`keyProviderID`).
2. Run the [RekeySoftwareEncryptionAtRestMasterKey](#) with the `keyManagementType` parameter as `external` and `keyProviderID` as the ID number of the key provider from the previous step:

```
{
  "method": "rekeysoftwareencryptionatrestmasterkey",
  "params": {
    "keyManagementType": "external",
    "keyProviderID": "<ID number>"
  }
}
```

3. Copy the `asyncHandle` value from the [RekeySoftwareEncryptionAtRestMasterKey](#) command response.
4. Run the [GetAsyncResult](#) command with the `asyncHandle` value from the previous step to confirm the change in configuration. From the command response, you should see that the older master key configuration has been updated with new key information. Copy the new key provider ID for use in a later step.

```

{
  "id": null,
  "result": {
    "createTime": "2021-01-01T22:29:18Z",
    "lastUpdateTime": "2021-01-01T22:45:51Z",
    "result": {
      "keyToDecommission": {
        "keyID": "<value>",
        "keyManagementType": "internal"
      },
      "newKey": {
        "keyID": "<value>",
        "keyManagementType": "external",
        "keyProviderID": <value>
      },
      "operation": "Rekeying Master Key. Master Key management being transferred from Internal Key Management to External Key Management with keyProviderID=<value>",
      "state": "Ready"
    },
    "resultType": "RekeySoftwareEncryptionAtRestMasterKey",
    "status": "complete"
  }
}

```

5. Run the `GetSoftwareEncryptionatRestInfo` command to confirm that new key details, including the `keyProviderID`, have been updated.

```

{
  "id": null,
  "result": {
    "masterKeyInfo": {
      "keyCreatedTime": "2021-01-01T22:29:18Z",
      "keyID": "<updated value>",
      "keyManagementType": "external",
      "keyProviderID": <value>
    },
    "rekeyMasterKeyAsyncResultID": <value>
  },
  "status": "enabled",
  "version": 1
}

```

**Find more information**

- [Manage storage with the Element API](#)
- [NetApp SolidFire Resources Page](#)
- [Documentation for earlier versions of NetApp SolidFire and Element products](#)

## Recover inaccessible or invalid authentication keys

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

### **The cluster is unable to unlock the drives due to a KmipServerFault cluster fault.**

This can occur when the cluster first boots up and the key server is inaccessible or the required key is unavailable.

1. Follow the recovery steps in the cluster fault codes (if any).

### **A sliceServiceUnhealthy fault might be set because the metadata drives have been marked as failed and placed into the "Available" state.**

Steps to clear:

1. Add the drives again.
2. After 3 to 4 minutes, check that the `sliceServiceUnhealthy` fault has cleared.

See [cluster fault codes](#) for more information.

## External key management API commands

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

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

- `CreatePublicPrivateKeyPair`
- `GetClientCertificateSignRequest`

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

- `CreateKeyServerKmip`
- `ModifyKeyServerKmip`
- `DeleteKeyServerKmip`
- `GetKeyServerKmip`
- `ListKeyServersKmip`
- `TestKeyServerKmip`

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

- `CreateKeyProviderKmip`
- `DeleteKeyProviderKmip`

- AddKeyServerToProviderKmip
- RemoveKeyServerFromProviderKmip
- GetKeyProviderKmip
- ListKeyProvidersKmip
- RekeySoftwareEncryptionAtRestMasterKey
- TestKeyProviderKmip

For information about the API methods, see [API reference information](#).

## 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.