



NetApp HCI 1.8 documentation

NetApp HCI

NetApp

June 25, 2025

This PDF was generated from <https://docs.netapp.com/us-en/hci18/index.html> on June 25, 2025. Always check docs.netapp.com for the latest.

Table of Contents

NetApp HCI 1.8 documentation	1
NetApp HCI Solutions	2
Release Notes	3
What's new in NetApp HCI	3
NetApp HCI 1.8P1	3
Element 12.2	3
Additional release information	5
NetApp HCI	5
NetApp Element software	5
Management services	5
NetApp Element Plug-in for vCenter Server	5
Compute firmware	6
Storage firmware	6
Concepts	7
NetApp HCI product overview	7
Components of NetApp HCI	7
NetApp HCI URLs	8
User accounts	8
User account management	8
Storage cluster administrator accounts	9
Authoritative user accounts	9
Volume accounts	10
Data protection	10
Remote replication types	10
Volume snapshots for data protection	12
Volume clones	12
Backup and restore process overview for SolidFire storage	12
Protection domains	13
Double Helix high availability	13
Clusters	13
Authoritative storage clusters	14
Stranded capacity	14
Two-node storage clusters	14
Storage clusters with three or more nodes	15
Find more information	16
Nodes	16
Management node	16
Storage nodes	16
Compute nodes	16
Witness Nodes	17
Find more information	17
Storage	17
Maintenance mode	17

Volumes	18
Volume access groups	19
Initiators	19
NetApp HCI licensing	20
NetApp HCI and VMware vSphere licensing	20
NetApp HCI and ONTAP Select licensing	20
Find more information	20
NetApp Hybrid Cloud Control configuration maximums	21
NetApp HCI security	21
Encryption at Rest for storage nodes	21
Software Encryption at Rest	22
External key management	22
Multi-factor authentication	22
FIPS 140-2 for HTTPS and data at rest encryption	22
Performance and Quality of Service	23
Quality of Service parameters	23
QoS value limits	24
QoS performance	24
QoS policies	25
Requirements and pre-deployment tasks	26
Requirements for NetApp HCI deployment overview	26
Network port requirements	26
Find more information	31
Network and switch requirements	31
Network cable requirements	32
IP address requirements	33
Number of IP addresses needed per NetApp HCI deployment	33
IP addresses reserved by NetApp HCI	33
Network configuration	34
Network configuration	34
Network configuration	35
Network configuration	37
Network configuration	39
DNS and timekeeping requirements	42
Environmental requirements	42
Protection domains	43
Witness Node resource requirements for two-node storage clusters	43
Get started with NetApp HCI	44
NetApp HCI installation and deployment overview	44
Prepare for installation	44
Validate network readiness with NetApp Active IQ Config Advisor	45
Work with your NetApp team	46
Install NetApp HCI hardware	46
Complete optional tasks after installing hardware	46
Deploy NetApp HCI using the NetApp Deployment Engine (NDE)	46

Manage NetApp HCI using the vCenter Plug-in	49
Monitor or upgrade NetApp HCI with the Hybrid Cloud Control	49
Install H-series hardware	50
Workflow diagrams	50
Prepare for installation	56
Install the rails	56
Install the node/chassis	58
Install the switches	60
Cable the nodes	60
Power on the nodes	64
Configure NetApp HCI	64
Perform post-configuration tasks	65
Configure LACP for optimal storage performance	66
Validate your environment with Active IQ Config Advisor	66
Install Active IQ Config Advisor	66
Use Active IQ Config Advisor	67
Configure IPMI for each node	69
Prerequisites	69
Manually assign the IPMI port IP address	69
Change the default IPMI password for H410C and H410S nodes	70
Change the default IPMI password for H610C, H615C, and H610S nodes	70
Deploy NetApp HCI	72
Access the NetApp Deployment Engine	72
Access the NetApp Deployment Engine	72
Access the NetApp Deployment Engine in environments without DHCP	72
Access the NetApp Deployment Engine in environments with DHCP	73
Manually assign IP addresses to access the NetApp Deployment Engine	74
Start your deployment	75
Configure VMware vSphere	75
VMware vSphere configuration	75
Configure a new VMware vSphere environment	76
Join an existing VMware vSphere deployment	77
Configuring NetApp HCI credentials	78
Select a network topology	78
Inventory selection	79
Inventory selection and node compatibility	79
Select inventory	80
Configure network settings	81
Find more information	84
Review and deploy the configuration	84
Find more information	85
Post-deployment tasks	86
Post-deployment tasks	86
Supported networking changes	86
Disable the smartd service on NetApp HCI compute nodes	88

Disable the "lacp-individual" command on configured switches	88
Keep VMware vSphere up to date	89
Install GPU drivers for GPU-enabled compute nodes	90
Configure Fully Qualified Domain Name web UI access	92
Access NetApp Hybrid Cloud Control	96
Reduce boot media wear on a NetApp HCI compute node	97
Manage NetApp HCI	98
NetApp HCI management overview	98
Update vCenter and ESXi credentials	98
Update vCenter password by using the management node REST API	98
Update the ESXi password by using the management node REST API	98
Manage NetApp HCI storage	100
Manage NetApp HCI storage overview	101
Create and manage user accounts by using NetApp Hybrid Cloud Control	101
Add and manage storage clusters using NetApp Hybrid Cloud Control	105
Create and manage volumes by using NetApp Hybrid Cloud Control	109
Create and manage volume access groups	115
Create and manage initiators	117
Create and manage volume QoS policies	119
Work with the management node	122
Management node overview	122
Install or recover a management node	123
Access the management node	137
Change the management node default SSL certificate	139
Work with the management node UI	140
Work with the management node REST API	145
Manage support connections	167
Power your NetApp HCI system off or on	170
Powering your NetApp HCI system off or on	170
Power off compute resources for a NetApp HCI system	171
Power off storage resources for a NetApp HCI system	172
Power on storage resources for a NetApp HCI system	173
Power on compute resources for a NetApp HCI system	173
Monitor your NetApp HCI system with NetApp Hybrid Cloud Control	175
Monitor storage and compute resources on the Hybrid Cloud Control Dashboard	175
Access the NetApp HCC Dashboard	175
Monitor storage resources	176
Monitor compute resources	177
Monitor storage capacity	177
Monitor storage performance	179
Monitor compute utilization	180
View your inventory on the Nodes page	181
Edit Baseboard Management Controller connection information	183
Use NetApp Hybrid Cloud Control to edit BMC information	183
Use the REST API to edit BMC information	184

Monitor volumes on your storage cluster	186
Monitor performance, capacity, and cluster health with SolidFire Active IQ	188
Collect logs for troubleshooting	189
Use NetApp Hybrid Cloud Control to collect logs	189
Use the REST API to collect logs	190
Upgrade your NetApp HCI system version 1.8	193
Upgrade sequence overview	193
System upgrade sequence	193
System upgrade procedures	194
Update management services	194
Upgrade to the latest HealthTools	197
Run Element storage health checks prior to upgrading storage	198
Upgrade Element software	207
Upgrade storage firmware	224
Upgrade a management node	233
Upgrade the Element Plug-in for vCenter Server	245
Run compute node health checks prior to upgrading compute firmware	249
Update compute node drivers	255
Upgrade compute node firmware	256
Upgrade your vSphere components for a NetApp HCI system with the Element Plug-in for vCenter Server	270
Expand your NetApp HCI system	271
Expansion overview	271
Expand NetApp HCI storage resources	271
Expand NetApp HCI compute resources	273
Expand NetApp HCI storage and compute resources at the same time	276
Remove Witness Nodes after expanding cluster	279
Find more information	280
Use Rancher on NetApp HCI	281
Rancher on NetApp HCI overview	281
Benefits of Rancher on NetApp HCI	281
What you can do with Rancher on NetApp HCI	281
Technical Support option	281
Rancher on NetApp HCI architecture and components	282
Rancher on NetApp HCI concepts	283
Trident software and persistent storage concepts	283
Requirements for Rancher on NetApp HCI	284
Node requirements	284
Network requirements	284
VMware vSphere requirements	285
Deployment considerations	285
Required ports	286
Required URLs	286
Deploy Rancher on NetApp HCI	286
What happens when you deploy Rancher on NetApp HCI?	287

Steps to deploy Rancher on NetApp HCI	287
What's next?	290
Post-deployment tasks	291
Post-deployment tasks overview	291
Ensure Rancher Support parity	291
Improve Rancher VM resiliency	291
Enable monitoring	292
Install Trident	292
Enable Trident support for user clusters	295
Deploy user clusters and applications	296
Deploy user clusters	296
Deploy applications on user clusters	296
Manage Rancher on NetApp HCI	297
Identify Rancher server cluster URLs and status	297
Monitor a Rancher on NetApp HCI implementation	297
Monitor Rancher using the NetApp Hybrid Cloud Control	298
Monitor Rancher using the Rancher UI	298
Monitor Rancher using NetApp Active IQ	298
Monitor Rancher using vCenter Server	298
Upgrade Rancher on NetApp HCI	299
Use NetApp Hybrid Cloud Control UI to upgrade a Rancher deployment	300
Use NetApp Hybrid Cloud Control API to upgrade a Rancher deployment	300
Remove a Rancher installation on NetApp HCI	305
Remove Rancher on NetApp HCI using NetApp Hybrid Cloud Control	305
Remove Rancher on NetApp HCI using the REST API	306
Maintain H-series hardware	308
H-series hardware maintenance overview	308
Replace 2U H-series chassis	308
Replace the chassis when additional unused space is available in the rack	309
Replace the chassis when additional unused space is not available in the rack	312
Find more information	314
Replace DC power supply units in H615C and H610S nodes	314
Find more information	317
Replace DIMMs in compute nodes	317
Prepare to replace the DIMM	317
Replace the DIMM from the chassis	320
Find more information	326
Replace drives for storage nodes	326
Best practices for handling drives	326
Best practices for adding and removing drives	327
Find more information	331
Replace H410C nodes	331
Prepare to replace the compute node	332
Replace the compute node in the chassis	333
Remove the compute node asset in NetApp HCI 1.7 and later	334

Add the compute node to the cluster	335
Redeploy Witness Nodes for two and three-node storage clusters	346
Change the password if you received a replacement node with a non-standard BMC password	349
Upgrade the BMC firmware on your node	349
Find more information	351
Replace H410S nodes	351
Prepare to replace the storage node	352
Replace the storage node in the chassis	353
Add the storage node to the cluster	354
Find more information	357
Replace H610C and H615C nodes	357
Prepare to replace the node	358
Replace the node	359
Add the node to the cluster	360
Install the GPU drivers	362
Find more information	363
Replace H610S nodes	364
Best practices for adding and removing drives	364
Find more information	366
Replace power supply units	366
Find more information	368
Replace SN2010, SN2100, and SN2700 switches	368
Prepare to replace the faulty switch	370
Create the configuration file	370
Remove the faulty switch and install the replacement	372
Verify the operating system version on the switch	373
Configure the replacement switch	374
Complete the replacement	375
Find more information	375
Replace storage node in a two-node cluster	376
Find more information	376
Legal notices	377
Copyright	377
Trademarks	377
Patents	377
Privacy policy	377
Open source	377

NetApp HCI 1.8 documentation

NetApp HCI Solutions

NetApp HCI enables you to deploy cloud services across multiple public cloud providers and on premise. You can use NetApp HCI to deploy services similar to what you might do with a cloud provider-all in a self-service mode without IT involvement.

To learn about NetApp HCI solutions, see the [NetApp HCI Solutions documentation](#).

Release Notes

What's new in NetApp HCI

NetApp periodically updates NetApp HCI to bring you new features, enhancements, and bug fixes. NetApp HCI 1.8P1 includes Element 12.2 for storage clusters.

- The [NetApp HCI 1.8P1](#) section describes new features and updates in NetApp HCI version 1.8P1.
- The [Element 12.2](#) section describes new features and updates in NetApp Element 12.2.

NetApp HCI 1.8P1

NetApp HCI 1.8P1 includes security and stability improvements.

NetApp HCI documentation enhancements

You can now access NetApp HCI upgrading, expansion, monitoring, and concepts information in an easy-to-navigate format [here](#).

NetApp Element Plug-in for vCenter Server 4.5 availability

The NetApp Element Plug-in for vCenter Server 4.5 is available outside of the management node 12.2 and NetApp HCI 1.8P1 releases. To upgrade the plug-in, follow the instructions in the [NetApp HCI Upgrades](#) documentation.

NetApp Hybrid Cloud Control enhancements

NetApp Hybrid Cloud Control is enhanced for version 1.8P1. [Learn more](#).

Element 12.2

NetApp HCI 1.8P1 includes Element 12.2 for storage clusters. Element 12.2 introduces SolidFire Enterprise SDS, software encryption at rest, maintenance mode, enhanced volume access security, Fully Qualified Domain Name (FQDN) access to UIs, storage node firmware updates, and security updates.

SolidFire Enterprise SDS

Element 12.2 introduces SolidFire Enterprise SDS (eSDS). SolidFire eSDS provides the benefits of SolidFire scale out technology and NetApp Element software data services on the hardware of your choice that meets the reference configuration for SolidFire eSDS. [Learn more](#).

The following are new Element API methods related to SolidFire eSDS ([Element 12.2 API information for SolidFire eSDS](#) has more information):

- `GetLicenseKey`
- `SetLicenseKey`

Software encryption at rest

Element 12.2 introduces software encryption at rest, which you can enable when you create a storage cluster (and is enabled by default when you create a SolidFire Enterprise SDS storage cluster). This feature encrypts

all data stored on the SSDs in the storage nodes and causes only a very small (~2%) performance impact on client IO.

The following are Element API methods related to software encryption at rest (the [Element API Reference Guide](#) has more information):

- CreateCluster

Maintenance mode

Element 12.2 introduces maintenance mode, which enables you to take a storage node offline for maintenance such as software upgrades or host repairs, while preventing a full sync of all data. If one or more nodes need maintenance, you can minimize the I/O impact to the rest of the storage cluster by enabling maintenance mode for those nodes before you begin. You can use maintenance mode with both appliance nodes as well as SolidFire eSDS nodes.

Enhanced volume access security

You can now restrict volume access to certain initiators based on VLAN (virtual network) association. You can associate new or existing initiators with one or more virtual networks, restricting that initiator to iSCSI targets accessible via those virtual networks.

The following are updated Element API methods related to these security improvements (the [Element API Reference Guide](#) has more information):

- CreateInitiators
- ModifyInitiators
- AddAccount
- ModifyAccount

Fully Qualified Domain Name (FQDN) access to UIs

Element 12.2 supports cluster web interface access using FQDNs. On Element 12.2 storage clusters, if you use the FQDN to access web user interfaces such as the Element web UI, per-node UI, or management node UI, you must first add a storage cluster setting to identify the FQDN used by the cluster. This setting enables the cluster to properly redirect a login session and facilitates better integration with external services like key managers and identity providers for multi-factor authentication. This feature requires management services version 2.15 or later. [Learn more](#).

Storage node firmware updates

Element 12.2 includes firmware updates for storage nodes. [Learn more](#).

Security enhancements

Element 12.2 resolves security vulnerabilities for storage nodes and the management node. [Learn more](#) about these security enhancements.

New SMART warning for failing drives

Element 12.2 now performs periodic health checks on SolidFire appliance drives using SMART health data from the drives. A drive that fails the SMART health check might be close to failure. If a drive fails the SMART

health check, a new critical severity cluster fault appears: Drive with serial: <serial number> in slot: <node slot><drive slot> has failed the SMART overall health check. To resolve this fault, replace the drive

Find more information

- [NetApp Hybrid Cloud Control and Management Services Release Notes](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)
- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)

Additional release information

You can find links to the latest and earlier release notes for various components of the NetApp HCI and Element storage environment.



You will be prompted to log in using your NetApp Support Site credentials.

NetApp HCI

- [NetApp HCI 1.8P1 Release Notes](#)
- [NetApp HCI 1.8 Release Notes](#)
- [NetApp HCI 1.7P1 Release Notes](#)

NetApp Element software

- [NetApp Element Software 12.2 Release Notes](#)
- [NetApp Element Software 12.0 Release Notes](#)
- [NetApp Element Software 11.8 Release Notes](#)
- [NetApp Element Software 11.7 Release Notes](#)
- [NetApp Element Software 11.5.1 Release Notes](#)
- [NetApp Element Software 11.3P1 Release Notes](#)

Management services

- [Management Services Release Notes](#)

NetApp Element Plug-in for vCenter Server

- [vCenter Plug-in 5.1 Release Notes](#) *NEW*
- [vCenter Plug-in 5.0 Release Notes](#)
- [vCenter Plug-in 4.10 Release Notes](#)
- [vCenter Plug-in 4.9 Release Notes](#)

- [vCenter Plug-in 4.8 Release Notes](#)
- [vCenter Plug-in 4.7 Release Notes](#)
- [vCenter Plug-in 4.6 Release Notes](#)
- [vCenter Plug-in 4.5 Release Notes](#)
- [vCenter Plug-in 4.4 Release Notes](#)
- [vCenter Plug-in 4.3 Release Notes](#)

Compute firmware

- [Compute Firmware Bundle 2.146 Release Notes](#)
- [Compute Firmware Bundle 2.27 Release Notes](#)
- [Compute Firmware Bundle 12.2.109 Release Notes](#)
- [Supported firmware and ESXi driver versions *NEW*](#)

Storage firmware

- [Storage Firmware Bundle 2.146 Release Notes](#)
- [Storage Firmware Bundle 2.99.2 Release Notes](#)
- [Storage firmware Bundle 2.76 Release Notes](#)
- [Storage firmware Bundle 2.27 Release Notes](#)
- [H610S BMC 3.84.07 Release Notes](#)
- [Supported firmware and ESXi driver versions *NEW*](#)

Concepts

NetApp HCI product overview

NetApp HCI is an enterprise-scale hybrid cloud infrastructure design that combines storage, compute, networking, and hypervisor—and adds capabilities that span public and private clouds.

NetApp's disaggregated hybrid cloud infrastructure allows independent scaling of compute and storage, adapting to workloads with guaranteed performance.

- Meets hybrid multicloud demand
- Scales compute and storage independently
- Simplifies data services orchestration across hybrid multiclouds

Components of NetApp HCI

Here is an overview of the various components of the NetApp HCI environment:

- NetApp HCI provides both storage and compute resources. You use the **NetApp Deployment Engine** wizard to deploy NetApp HCI. After successful deployment, compute nodes appear as ESXi hosts and you can manage them in VMware vSphere Web Client.
- **Management services** or microservices include the Active IQ collector, QoSIOC for the vCenter Plug-in, and mNode service; they are updated frequently as service bundles. As of the Element 11.3 release, **management services** are hosted on the management node, allowing for quicker updates of select software services outside of major releases. The **management node** (mNode) is a virtual machine that runs in parallel with one or more Element software-based storage clusters. It is used to upgrade and provide system services including monitoring and telemetry, manage cluster assets and settings, run system tests and utilities, and enable NetApp Support access for troubleshooting.



Learn more about [management services releases](#).

- **NetApp Hybrid Cloud Control** enables you to manage NetApp HCI. You can upgrade management services, expand your system, collect logs, and monitor your installation by using NetApp SolidFire Active IQ. You log in to NetApp Hybrid Cloud Control by browsing to the IP address of the management node.
- The **NetApp Element Plug-in for vCenter Server** is a web-based tool integrated with the vSphere user interface (UI). The plug-in is an extension and scalable, user-friendly interface for VMware vSphere that can manage and monitor storage clusters running **NetApp Element software**. The plug-in provides an alternative to the Element UI. 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.



Learn more about [NetApp Element Plug-in for vCenter Server](#).

- By default, NetApp HCI sends performance and alert statistics to the **NetApp SolidFire Active IQ** service.

As part of your normal support contract, NetApp Support monitors this data and alerts you to any performance bottlenecks or potential system issues. You need to create a NetApp Support account if you do not already have one (even if you have an existing SolidFire Active IQ account) so that you can take advantage of this service.



Learn more about [NetApp SolidFire Active IQ](#).

NetApp HCI URLs

Here are the common URLs you use with NetApp HCI:

URL	Description
<code>https://[IPv4 address of Bond1G interface on a storage node]</code>	Access the NetApp Deployment Engine wizard to install and configure NetApp HCI. Learn more .
<code><code>https://&lt;ManagementNodeIP&gt;;</code></code>	Access NetApp Hybrid Cloud Control to upgrade, expand, and monitor your NetApp HCI installation, and update management services. Learn more .
<code>https://[IP address]:442</code>	From the per-node UI, access network and cluster settings and utilize system tests and utilities. Learn more .
<code>https://<ManagementNodeIP>:9443</code>	Register the vCenter Plug-in package in the vSphere Web Client.
<code>https://activeiq.solidfire.com</code>	Monitor data and receive alerts to any performance bottlenecks or potential system issues.
<code>https://<ManagementNodeIP>/mnode</code>	Manually update management services using the REST API UI from the management node.
<code>https://[storage cluster MVIP address]</code>	Access the NetApp Element software UI.

Find more information

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

User accounts

To access storage resources on your system, you'll need to set up user accounts.

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.

Here are some additional considerations:

- The account contains the CHAP authentication required to access the volumes assigned to it.
- An account can have up to 2000 volumes assigned to it, but a volume can belong to only one account.
- User accounts can be managed from the NetApp Element Management extension point.

Using NetApp Hybrid Cloud Control, you can create and manage the following types of accounts:

- Administrator user accounts for the storage cluster
- Authoritative user accounts
- Volume accounts, specific only to the storage cluster on which they were created.

Storage cluster administrator accounts

There are two types of administrator accounts that can exist in a storage cluster running NetApp Element software:

- **Primary cluster administrator account:** This administrator account is created when the cluster is created. This account is the primary administrative account with the highest level of access to the cluster. This account is analogous to a root user in a Linux system. You can change the password for this administrator account.
- **Cluster administrator account:** You can give a cluster administrator account a limited range of administrative access to perform specific tasks within a cluster. The credentials assigned to each cluster administrator account are used to authenticate API and Element UI requests within the storage system.



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

You can manage cluster administrator accounts by creating, deleting, and editing cluster administrator accounts, changing the cluster administrator password, and configuring LDAP settings to manage system access for users.

For details, see the [SolidFire and Element Documentation Center](#).

Authoritative user accounts

Authoritative user accounts can authenticate against any storage asset associated with the NetApp Hybrid Cloud Control instance of nodes and clusters. With this account, you can manage volumes, accounts, access groups, and more across all clusters.

Authoritative user accounts are managed from the top right menu User Management option in NetApp Hybrid Cloud Control.

The [authoritative storage cluster](#) is the storage cluster that NetApp Hybrid Cloud Control uses to authenticate users.

All users created on the authoritative storage cluster can log into the NetApp Hybrid Cloud Control. Users created on other storage clusters *cannot* log into Hybrid Cloud Control.

- If your management node only has one storage cluster, then it is the authoritative cluster.
- If your management node has two or more storage clusters, one of those clusters is assigned as the authoritative cluster and only users from that cluster can log into NetApp Hybrid Cloud Control.

While many NetApp Hybrid Cloud Control features work with multiple storage clusters, authentication and authorization have necessary limitations. The limitation around authentication and authorization is that users from the authoritative cluster can execute actions on other clusters tied to NetApp Hybrid Cloud Control even if they are not a user on the other storage clusters. Before proceeding with managing multiple storage clusters, you should ensure that users defined on the authoritative clusters are defined on all other storage clusters with the same permissions. You can manage users from NetApp Hybrid Cloud Control.

Volume accounts

Volume-specific accounts are specific only to the storage cluster on which they were created. These accounts enable you to set permissions on specific volumes across the network, but have no effect outside of those volumes.

Volume accounts are managed within the NetApp Hybrid Cloud Control Volumes table.

Find more information

- [Manage user accounts](#)
- [Learn about clusters](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [SolidFire and Element Documentation Center](#)

Data protection

NetApp HCI data protection terms include different types of remote replication, volume snapshots, volume cloning, protection domains, and high availability with double Helix technology.

NetApp HCI data protection includes the following concepts:

- [Remote replication types](#)
- [Volume snapshots for data protection](#)
- [Volume clones](#)
- [Backup and restore process overview for SolidFire storage](#)
- [Protection domains](#)
- [Double Helix high availability](#)

Remote replication types

Remote replication of data can take the following forms:

- [Synchronous and asynchronous replication between clusters](#)
- [Snapshot-only replication](#)
- [Replication between Element and ONTAP clusters using SnapMirror](#)

See [TR-4741: NetApp Element Software Remote Replication](#).

Synchronous and asynchronous replication between clusters

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 fallback scenarios.

Synchronous replication

Synchronous replication continuously replicates data from the source cluster to the target cluster and is affected by latency, packet loss, jitter, and bandwidth.

Synchronous replication is appropriate for the following situations:

- Replication of several systems over a short distance
- A disaster recovery site that is geographically local to the source
- Time-sensitive applications and the protection of databases
- Business continuity applications that require the secondary site to act as the primary site when the primary site is down

Asynchronous replication

Asynchronous replication continuously replicates data from a source cluster to a target cluster without waiting for the acknowledgments from the target cluster. During asynchronous replication, writes are acknowledged to the client (application) after they are committed on the source cluster.

Asynchronous replication is appropriate for the following situations:

- The disaster recovery site is far from the source and the application does not tolerate latencies induced by the network.
- There are bandwidth limitations on the network connecting the source and target clusters.

Snapshot-only replication

Snapshot-only data protection replicates changed data at specific points of time to a remote cluster. Only those snapshots that are created on the source cluster are replicated. Active writes from the source volume are not.

You can set the frequency of the snapshot replications.

Snapshot replication does not affect asynchronous or synchronous replication.

Replication between Element and ONTAP clusters using SnapMirror

With NetApp SnapMirror technology, you can replicate snapshots that were taken using NetApp Element software to ONTAP for disaster recovery purposes. In a SnapMirror relationship, Element is one endpoint and ONTAP is the other.

SnapMirror is a NetApp Snapshot™ replication technology that facilitates disaster recovery, designed for failover from primary storage to secondary storage at a geographically remote site. SnapMirror technology creates a replica, or mirror, of the working data in secondary storage from which you can continue to serve data if an outage occurs at the primary site. Data is mirrored at the volume level.

The relationship between the source volume in primary storage and the destination volume in secondary

storage is called a data protection relationship. The clusters are referred to as endpoints in which the volumes reside and the volumes that contain the replicated data must be peered. A peer relationship enables clusters and volumes to exchange data securely.

SnapMirror runs natively on the NetApp ONTAP controllers and is integrated into Element, which runs on NetApp HCI and SolidFire clusters. The logic to control SnapMirror resides in ONTAP software; therefore, all SnapMirror relationships must involve at least one ONTAP system to perform the coordination work. Users manage relationships between Element and ONTAP clusters primarily through the Element UI; however, some management tasks reside in NetApp ONTAP System Manager. Users can also manage SnapMirror through the CLI and API, which are both available in ONTAP and Element.

See [TR-4651: NetApp SolidFire SnapMirror Architecture and Configuration](#) (login required).

You must manually enable SnapMirror functionality at the cluster level by using Element software. SnapMirror functionality is disabled by default, and it is not automatically enabled as part of a new installation or upgrade.

After enabling SnapMirror, you can create SnapMirror relationships from the Data Protection tab in the Element software.

Volume snapshots for data protection

A volume snapshot is a point-in-time copy of a volume that you could later use to restore a volume to that specific time.

While snapshots are similar to volume clones, snapshots are simply replicas of volume metadata, so you cannot mount or write to them. Creating a volume snapshot also takes only a small amount of system resources and space, which makes snapshot creation faster than cloning.

You can replicate snapshots to a remote cluster and use them as a backup copy of the volume. This enables you to roll back a volume to a specific point in time by using the replicated snapshot; you can also create a clone of a volume from a replicated snapshot.

You can back up snapshots from a SolidFire cluster to an external object store, or to another SolidFire cluster. When you back up a snapshot to an external object store, you must have a connection to the object store that allows read/write operations.

You can take a snapshot of an individual volume or multiple for data protection.

Volume clones

A clone of a single volume or multiple volumes is 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.

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

Backup and restore process overview for SolidFire storage

You can back up and restore volumes to other SolidFire storage, as well as to secondary object stores that are compatible with Amazon S3 or OpenStack Swift.

You can back up a volume to the following:

- A SolidFire storage cluster
- An Amazon S3 object store
- An OpenStack Swift object store

When you restore volumes from OpenStack Swift or Amazon S3, you need manifest information from the original backup process. If you are restoring a volume that was backed up on a SolidFire storage system, no manifest information is required.

Protection domains

A protection domain is a node or a set of nodes grouped together such that any part or even all of it might fail, while maintaining data availability. Protection domains enable a storage cluster to heal automatically from the loss of a chassis (chassis affinity) or an entire domain (group of chassis).

A protection domain layout assigns each node to a specific protection domain.

Two different protection domain layouts, called protection domain levels, are supported.

- At the node level, each node is in its own protection domain.
- At the chassis level, only nodes that share a chassis are in the same protection domain.
 - The chassis level layout is automatically determined from the hardware when the node is added to the cluster.
 - In a cluster where each node is in a separate chassis, these two levels are functionally identical.

You can manually [enable protection domain monitoring](#) using the NetApp Element Plug-in for vCenter Server. You can select a protection domain threshold based on node or chassis domains.

When creating a new cluster, if you are using storage nodes that reside in a shared chassis, you might want to consider designing for chassis-level failure protection using the protection domains feature.

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

See [SolidFire and Element 12.2 Documentation Center](#).

Double Helix high availability

Double Helix data protection is a replication method that spreads at least two redundant copies of data across all drives within a system. The “RAID-less” approach enables a system to absorb multiple, concurrent failures across all levels of the storage system and repair quickly.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Clusters

A cluster is a group of nodes, functioning as a collective whole, that provide storage or compute resources. Starting with NetApp HCI 1.8, you can have a storage cluster with two nodes. A storage cluster appears on the network as a single logical group and can

then be accessed as block storage.

The storage layer in NetApp HCI is provided by NetApp Element software and the management layer is provided by the NetApp Element Plug-in for vCenter Server. A storage node is a server containing a collection of drives that communicate with each other through the Bond10G network interface. Each storage node is connected to two networks, storage and management, each with two independent links for redundancy and performance. Each node requires an IP address on each network. You can create a cluster with new storage nodes, or add storage nodes to an existing cluster to increase storage capacity and performance.

Authoritative storage clusters

The authoritative storage cluster is the storage cluster that NetApp Hybrid Cloud Control uses to authenticate users.

If your management node only has one storage cluster, then it is the authoritative cluster. If your management node has two or more storage clusters, one of those clusters is assigned as the authoritative cluster and only users from that cluster can log into NetApp Hybrid Cloud Control. To find out which cluster is the authoritative cluster, you can use the `GET /mnode/about` API. In the response, the IP address in the `token_url` field is the management virtual IP address (MVIP) of the authoritative storage cluster. If you attempt to log into NetApp Hybrid Cloud Control as a user that is not on the authoritative cluster, the login attempt will fail.

Many NetApp Hybrid Cloud Control features are designed to work with multiple storage clusters, but authentication and authorization have limitations. The limitation around authentication and authorization is that the user from the authoritative cluster can execute actions on other clusters tied to NetApp Hybrid Cloud Control even if they are not a user on the other storage clusters. Before proceeding with managing multiple storage clusters, you should ensure that users defined on the authoritative clusters are defined on all other storage clusters with the same permissions.

You can manage users with NetApp Hybrid Cloud Control.

Before proceeding with managing multiple storage clusters, you should ensure that users defined on the authoritative clusters are defined on all other storage clusters with the same permissions. You can [manage users](#) from the Element software user interface (Element web UI).

See [Create and manage storage cluster assets](#) for more information on working with management node storage cluster assets.

Stranded capacity

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 capacity 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 from happening. When a node becomes stranded, an appropriate cluster fault is thrown.

Two-node storage clusters

Starting with NetApp HCI 1.8, you can set up a storage cluster with two storage nodes.

- You can use certain types of nodes to form the two-node storage cluster. See [NetApp HCI 1.8 Release Notes](#).



In a two-node cluster, the storage nodes are limited to nodes with 480GB and 960GB drives and the nodes must be the same model type.

- Two-node storage clusters are best suited for small-scale deployments with workloads that are not dependent on large capacity and high performance requirements.
- In addition to two storage nodes, a two-node storage cluster also includes two **NetApp HCI Witness Nodes**.



Learn more about [Witness Nodes](#).

- You can scale a two-node storage cluster to a three-node storage cluster. Three-node clusters increase resiliency by providing the ability to auto-heal from storage node failures.
- Two-node storage clusters provide the same security features and functionality as the traditional four-node storage clusters.
- Two-node storage clusters use the same networks as four-node storage clusters. The networks are set up during NetApp HCI deployment using the NetApp Deployment Engine wizard.

Storage cluster quorum

Element software creates a storage cluster from selected nodes, which maintains a replicated database of the cluster configuration. A minimum of three nodes are required to participate in the cluster ensemble to maintain quorum for cluster resiliency. Witness Nodes in a two-node cluster are used to ensure that there are enough storage nodes to form a valid ensemble quorum. For ensemble creation, storage nodes are preferred over Witness Nodes. For the minimum three-node ensemble involving a two-node storage cluster, two storage nodes and one Witness Node are used.



In a three-node ensemble with two storage nodes and one Witness Node, if one storage node goes offline, the cluster goes into a degraded state. Of the two Witness Nodes, only one can be active in the ensemble. The second Witness Node cannot be added to the ensemble, because it performs the backup role. The cluster stays in degraded state until the offline storage node returns to an online state, or a replacement node joins the cluster.

If a Witness Node fails, the remaining Witness Node joins the ensemble to form a three-node ensemble. You can deploy a new Witness Node to replace the failed Witness Node.

Auto-healing and failure handling in two-node storage clusters

If a hardware component fails in a node that is part of a traditional cluster, the cluster can rebalance data that was on the component that failed to other available nodes in the cluster. This ability to automatically heal is not available in a two-node storage cluster, because a minimum of three physical storage nodes must be available to the cluster for healing automatically. When one node in a two-node cluster fails, the two-node cluster does not require regeneration of a second copy of data. New writes are replicated for block data in the remaining active storage node. When the failed node is replaced and joins the cluster, the data is rebalanced between the two physical storage nodes.

Storage clusters with three or more nodes

Expanding from two storage nodes to three storage nodes makes your cluster more resilient by allowing auto-healing in the event of node and drive failures, but does not provide additional capacity. You can expand using the [NetApp Hybrid Cloud Control UI](#). When expanding from a two-node cluster to a three-node cluster, capacity can be stranded (see [Stranded capacity](#)). The UI wizard shows warnings about stranded capacity before

installation. A single Witness Node is still available to keep the ensemble quorum in the event of a storage node failure, with a second Witness Node on standby.

When you expand a three-node storage cluster to a four-node cluster, capacity and performance are increased. In a four-node cluster, Witness Nodes are no longer needed to form the cluster quorum. You can expand to up to 64 compute nodes and 40 storage nodes.

Find more information

- [NetApp HCI Two-Node Storage Cluster | TR-4823](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [SolidFire and Element Software Documentation Center](#)

Nodes

Nodes are hardware or virtual resources that are grouped into a cluster to provide block storage and compute capabilities.

NetApp HCI and Element software defines various node roles for a cluster. The four types of node roles are **management node**, **storage node**, **compute node**, and **NetApp HCI Witness Nodes**.

Management node

The management node (sometimes abbreviated as mNode) interacts with a storage cluster to perform management actions, but is not a member of the storage cluster. Management nodes periodically collect information about the cluster through API calls and report this information to Active IQ for remote monitoring (if enabled). Management nodes are also responsible for coordinating software upgrades of the cluster nodes.

The management node is a virtual machine (VM) that runs in parallel with one or more Element software-based storage clusters. In addition to upgrades, it is used to provide system services including monitoring and telemetry, manage cluster assets and settings, run system tests and utilities, and enable NetApp Support access for troubleshooting. As of the Element 11.3 release, the management node functions as a microservice host, allowing for quicker updates of select software services outside of major releases. These microservices or management services, such as the Active IQ collector, QoSIOC for the vCenter Plug-in, and management node service, are updated frequently as service bundles.

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.

Compute nodes

NetApp HCI compute nodes are hardware that provides compute 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. Regardless of whether it is a four-node storage cluster or a two-node storage cluster, the minimum number of compute nodes remains two for a NetApp HCI deployment.

Witness Nodes

NetApp HCI Witness Nodes are VMs that run on compute nodes in parallel with an Element software-based storage cluster. Witness Nodes do not host slice or block services. A Witness Node enables storage cluster availability in the event of a storage node failure. You can manage and upgrade Witness Nodes in the same way as other storage nodes. A storage cluster can have up to four Witness Nodes. Their primary purpose is to ensure that enough cluster nodes exist to form a valid ensemble quorum.

Best practice: Configure the Witness Node VMs to use the compute node's local datastore (default set by NDE), do not configure them on shared storage, such as SolidFire storage volumes. To prevent the VMs migrating automatically, set the Witness Node VM's Distributed Resource Scheduler (DRS) automation level to **Disabled**. This prevents both Witness Nodes running on the same compute node and creating a non-high availability (HA) pair configuration.



Learn more about [Witness Node resource requirements](#) and [Witness Node IP address requirements](#).



In a two-node storage cluster, a minimum of two Witness Nodes are deployed for redundancy in the event of a Witness Node failure. When the NetApp HCI installation process installs Witness Nodes, a VM template is stored in VMware vCenter that you can use to redeploy a Witness Node in case it is accidentally removed, lost, or corrupted. You can also use the template to redeploy a Witness Node if you need to replace a failed compute node that was hosting the Witness Node. For instructions, see the **Redeploy Witness Nodes for two and three-node storage clusters** section [here](#).

Find more information

- [NetApp HCI Two-Node Storage Cluster | TR-4823](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [SolidFire and Element Software Documentation Center](#)

Storage

Maintenance mode

If you need to take a storage node offline for maintenance such as software upgrades or host repairs, you can minimize the I/O impact to the rest of the storage cluster by enabling maintenance mode for that node. You can use maintenance mode with both appliance nodes as well as SolidFire Enterprise SDS nodes.

You can only transition a storage node to maintenance mode if the node is healthy (has no blocking cluster faults) and the storage cluster is tolerant to a single node failure. Once you enable maintenance mode for a healthy and tolerant node, the node is not immediately transitioned; it is monitored until the following conditions are true:

- All volumes hosted on the node have failed over
- The node is no longer hosting as the primary for any volume
- A temporary standby node is assigned for every volume being failed over

After these criteria are met, the node is transitioned to maintenance mode. If these criteria are not met within a 5 minute period, the node will not enter maintenance mode.

When you disable maintenance mode for a storage node, the node is monitored until the following conditions are true:

- All data is fully replicated to the node
- All blocking cluster faults are resolved
- All temporary standby node assignments for the volumes hosted on the node have been inactivated

After these criteria are met, the node is transitioned out of maintenance mode. If these criteria are not met within one hour, the node will fail to transition out of maintenance mode.

You can see the states of maintenance mode operations when working with maintenance mode using the Element API:

- **Disabled:** No maintenance has been requested.
- **FailedToRecover:** The node failed to recover from maintenance.
- **RecoveringFromMaintenance:** The node is in the process of recovering from maintenance.
- **PreparingForMaintenance:** Actions are being taken to allow a node to have maintenance performed.
- **ReadyForMaintenance:** The node is ready for maintenance to be performed.

Find more information

- [SolidFire and Element Documentation Center](#)

Volumes

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.

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 yet recommended management node configuration.

If you are deploying a management node for NetApp HCI using the NetApp Deployment Engine, persistent volumes are enabled and configured automatically.

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.



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

Find more information

- [Manage volumes](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [SolidFire and Element Software Documentation Center](#)

Volume access groups

A volume access group is a collection of volumes that users can access using either iSCSI or Fibre Channel initiators.

By creating and using volume access groups, you can control access to a set of volumes. When you associate a set of volumes and a set of initiators with a volume access group, the access group grants those initiators access to that set of volumes.

Volume access groups have the following limits:

- A maximum of 128 initiators per volume access group.
- A maximum of 64 access groups per volume.
- An access group can be made up of a maximum of 2000 volumes.
- An IQN or WWPN can belong to only one volume access group.

Find more information

- [Manage volume access groups](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [SolidFire and Element Software Documentation Center](#)

Initiators

Initiators enable external clients access to volumes in a cluster, serving as the entry point for communication between clients and volumes. You can use initiators for CHAP-based rather than account-based access to storage volumes. A single initiator, when added to a volume access group, allows volume access group members to access all storage volumes added to the group without requiring authentication. An initiator can belong to only one access group.

Find more information

- [Manage initiators](#)
- [Volume access groups](#)
- [Manage volume access groups](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [SolidFire and Element Software Documentation Center](#)

NetApp HCI licensing

When you use NetApp HCI, you might need additional licenses depending on what you are using.

NetApp HCI and VMware vSphere licensing

VMware vSphere licensing depends on your configuration:

Networking option	Licensing
Option A: Two cables for compute nodes using VLAN tagging (All compute nodes)	Requires use of vSphere Distributed Switch, which requires VMware vSphere Enterprise Plus licensing.
Option B: Six cables for compute nodes using tagged VLANs (H410C 2RU 4-Node compute node)	This configuration uses vSphere Standard Switch as the default. Optional use of vSphere Distributed Switch requires VMware Enterprise Plus licensing.
Option C: Six cables for compute nodes using native and tagged VLANs (H410C, 2RU 4-Node Compute Node)	This configuration uses vSphere Standard Switch as the default. Optional use of vSphere Distributed Switch requires VMware Enterprise Plus licensing.

NetApp HCI and ONTAP Select licensing

If you were provided a version of ONTAP Select for use in conjunction with a purchased NetApp HCI system, the following additional limitations apply:

- The ONTAP Select license, which is bundled with a NetApp HCI system sale, may only be used in conjunction with NetApp HCI compute nodes.
- The storage for those ONTAP Select instances must reside only on the NetApp HCI storage nodes.
- The use of third-party compute nodes or third-party storage nodes is prohibited.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [SolidFire and Element Software Documentation Center](#)

NetApp Hybrid Cloud Control configuration maximums

NetApp HCI includes NetApp Hybrid Cloud Control to simplify compute lifecycle and storage management. It supports Element Software upgrades on storage nodes for NetApp HCI and NetApp SolidFire storage clusters, as well as firmware upgrades for NetApp HCI compute nodes in NetApp HCI. It is available by default on the management nodes in NetApp HCI.

In addition to communicating the NetApp-provided hardware and software components in a NetApp HCI installation, NetApp Hybrid Cloud Control interacts with third-party components in the customer environment, like VMware vCenter. NetApp qualifies the functionality of NetApp Hybrid Cloud Control and its interaction with these third-party components in the customer environment up to a certain scale. For optimal experience with NetApp Hybrid Cloud Control, NetApp recommends staying within the range of configuration maximums.

If you exceed these tested maximums, you might experience issues with NetApp Hybrid Cloud Control, such as a slower user interface and API responses or functionality being unavailable. If you engage NetApp for product support with NetApp Hybrid Cloud Control in environments that are configured beyond the configuration maximums, NetApp Support will ask that you change the configuration to be within the documented configuration maximums.

Configuration maximums

NetApp Hybrid Cloud Control supports VMware vSphere environments with up to 100 ESXi hosts and 1000 virtual machines (comparable to a small vCenter Server Appliance configuration).

NetApp HCI security

When you use NetApp HCI, your data is protected by industry-standard security protocols.

Encryption at Rest for storage nodes

NetApp HCI enables you to encrypt all data stored on the storage cluster.

All drives in storage nodes that are capable of encryption use AES 256-bit encryption at the drive level. Each drive has its own encryption key, which is created when the drive is first initialized. When you enable the encryption feature, a storage-cluster-wide password is created, and chunks of the password are then distributed to all nodes in the cluster. No single node stores the entire password. The password is then used to password-protect all access to the drives. You need the password to unlock the drive, and since the drive is encrypting all data, your data is secure at all times.

When you enable Encryption at Rest, performance and efficiency of the storage cluster are unaffected. Additionally, if you remove an encryption-enabled drive or node from the storage cluster with the Element API or Element UI, Encryption at Rest is disabled on the drives and the drives are securely erased, protecting the data that was previously stored on those drives. After you remove the drive, you can securely erase the drive with the `SecureEraseDrives` API method. If you forcibly remove a drive or node from the storage cluster, the data remains protected by the cluster-wide password and the drive's individual encryption keys.

For information on enabling and disabling Encryption at Rest, see [Enabling and disabling encryption for a cluster](#) in the SolidFire and Element Documentation Center.

Software Encryption at Rest

Software Encryption at Rest enables all data written to the SSDs in a storage cluster to be encrypted. This provides a primary layer of encryption in SolidFire Enterprise SDS nodes that do not include Self-Encrypting Drives (SEDs).

External key management

You can configure Element software to use a third-party KMIP-compliant key management service (KMS) to manage storage cluster encryption keys. When you enable this feature, the storage cluster's cluster-wide drive access password encryption key is managed by a KMS that you specify.

Element can use the following key management services:

- Gemalto SafeNet KeySecure
- SafeNet AT KeySecure
- HyTrust KeyControl
- Vormetric Data Security Manager
- IBM Security Key Lifecycle Manager

For more information on configuring External Key Management, see [Getting started with External Key Management](#) in the SolidFire and Element Documentation Center.

Multi-factor authentication

Multi-factor authentication (MFA) enables you to require users to present multiple types of evidence to authenticate with the NetApp Element web UI or storage node UI upon login. You can configure Element to accept only multi-factor authentication for logins integrating with your existing user management system and identity provider.

You can configure Element to integrate with an existing SAML 2.0 identity provider which can enforce multiple authentication schemes, such as password and text message, password and email message, or other methods.

You can pair multi-factor authentication with common SAML 2.0 compatible identity providers (IdPs), such as Microsoft Active Directory Federation Services (ADFS) and Shibboleth.

To configure MFA, see [Enabling multi-factor authentication](#) in the SolidFire and Element Documentation Center.

FIPS 140-2 for HTTPS and data at rest encryption

NetApp SolidFire storage clusters and NetApp HCI systems support encryption that complies with the Federal Information Processing Standard (FIPS) 140-2 requirements for cryptographic modules. You can enable FIPS 140-2 compliance on your NetApp HCI or SolidFire cluster for both HTTPS communications and drive encryption.

When you enable FIPS 140-2 operating mode on your cluster, the cluster activates the NetApp Cryptographic Security Module (NCSM) and leverages FIPS 140-2 Level 1 certified encryption for all communication via HTTPS to the NetApp Element UI and API. You use the `EnableFeature` Element API with the `fips` parameter to enable FIPS 140-2 HTTPS encryption. On storage clusters with FIPS-compatible hardware, you can also enable FIPS drive encryption for data at rest using the `EnableFeature` Element API with the `FipsDrives` parameter.

For more information about preparing a new storage cluster for FIPS 140-2 encryption, see [Creating a cluster supporting FIPS drives](#).

For more information about enabling FIPS 140-2 on an existing, prepared cluster, see [The EnableFeature Element API](#).

Performance and Quality of Service

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



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

Quality of Service parameters

IOPS parameters are defined in the following ways:

- **Minimum IOPS** - The minimum number of sustained inputs and outputs per second (IOPS) that the storage cluster provides to a volume. The Min IOPS configured for a volume is the guaranteed level of performance for a volume. Performance does not drop below this level.
- **Maximum IOPS** - The maximum number of sustained IOPS that the storage cluster provides to a volume. When cluster IOPS levels are critically high, this level of IOPS performance is not exceeded.
- **Burst IOPS** - The maximum number of IOPS allowed in a short burst scenario. If a volume has been running below the Max IOPS, burst credits are accumulated. When performance levels become very high and are pushed to maximum levels, short bursts of IOPS are allowed on the volume.

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

A single volume can accrue Burst IOPS and use the credits to burst above their Max IOPS up to their Burst IOPS level for a set "burst period." A volume can burst for up to 60 seconds if the cluster has the capacity to accommodate the burst. A volume accrues one second of burst credit (up to a maximum of 60 seconds) for every second that the volume runs below its Max IOPS limit.

Burst IOPS are limited in two ways:

- A volume can burst above its Max IOPS for a number of seconds equal to the number of burst credits that the volume has accrued.
 - When a volume bursts above its Max IOPS setting, it is limited by its Burst IOPS setting. Therefore, the burst IOPS never exceeds the burst IOPS setting for the volume.
- **Effective Max Bandwidth** - The maximum bandwidth is calculated by multiplying the number of IOPS (based on the QoS curve) by the IO size.

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

- Workloads are able to reach and sustain a maximum of 1000 IOPS until the condition of workload contention for IOPS becomes apparent on the cluster. IOPS are then reduced incrementally until IOPS on all volumes are within the designated QoS ranges and contention for performance is relieved.

- Performance on all volumes is pushed toward the Min IOPS of 100. Levels do not drop below the Min IOPS setting but could remain higher than 100 IOPS when workload contention is relieved.
- Performance is never greater than 1000 IOPS, or less than 100 IOPS for a sustained period. Performance of 1500 IOPS (Burst IOPS) is allowed, but only for those volumes that have accrued burst credits by running below Max IOPS and only allowed for a short periods of time. Burst levels are never sustained.

QoS value limits

Here are the possible minimum and maximum values for QoS.

Parameters	Min value	Default	4 4KB	5 8KB	6 16KB	262KB
Min IOPS	50	50	15,000	9,375*	5556*	385*
Max IOPS	100	15,000	200,000**	125,000	74,074	5128
Burst IOPS	100	15,000	200,000**	125,000	74.074	5128

*These estimations are approximate.

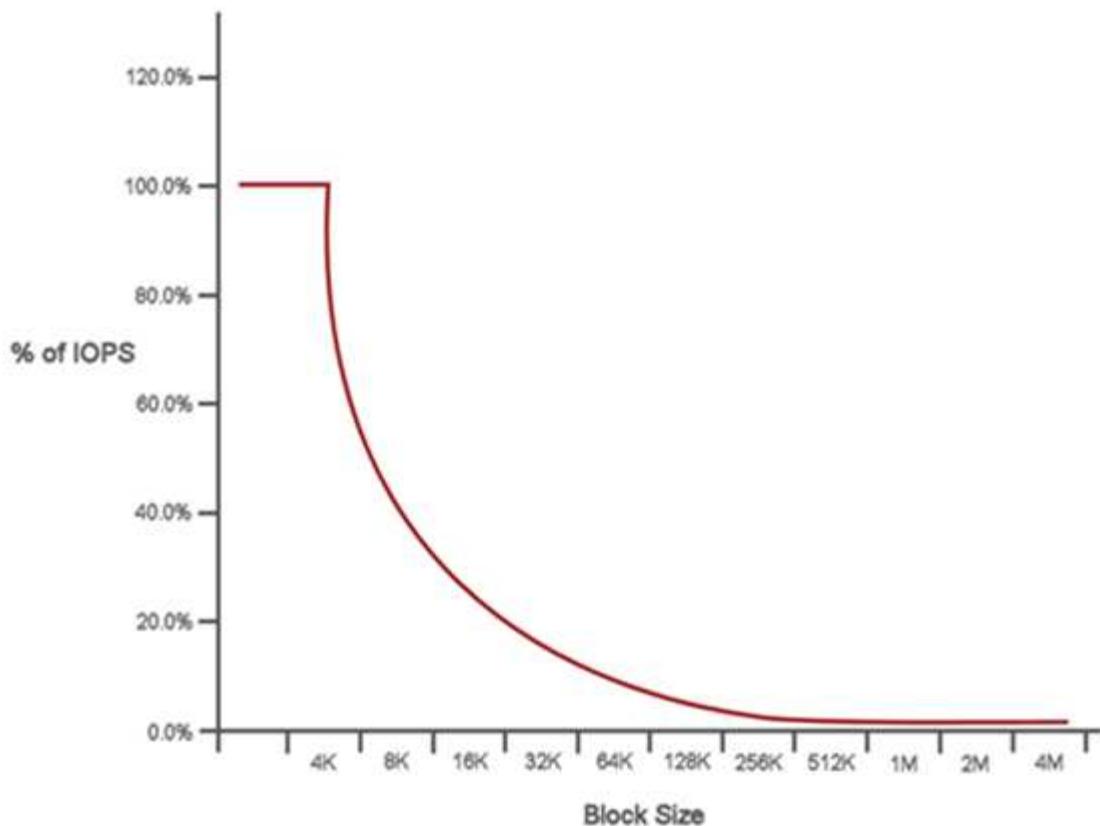
**Max IOPS and Burst IOPS can be set as high as 200,000; however, this setting is allowed only to effectively uncap the performance of a volume. Real-world maximum performance of a volume is limited by cluster usage and per-node performance.

QoS performance

The QoS performance curve shows the relationship between block size and the percentage of IOPS.

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

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



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

QoS policies

A QoS policy enables you to create and save a standardized quality of service setting that can be applied to many volumes.

QoS policies are best for service environments, for example, with database, application, or infrastructure servers that rarely reboot and need constant equal access to storage. Individual volume QoS is best for light use VMs, such as virtual desktops or specialized kiosk-type VMs, that may be rebooted, powered on, or powered off daily or several times a day.

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



The selected cluster must be Element 10.0 or later to use QoS policies; otherwise, QoS policy functions are not available.

Find more information

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

Requirements and pre-deployment tasks

Requirements for NetApp HCI deployment overview

NetApp HCI has specific physical and network requirements for proper operation in your datacenter. Ensure that you implement the following requirements and recommendations before you begin deployment.

Before you receive your NetApp HCI hardware, ensure that you complete the checklist items in the pre-deployment workbook from NetApp Professional Services. This document contains a comprehensive list of tasks you need to complete to prepare your network and environment for a successful NetApp HCI deployment.

Here are the links to the requirements and pre-deployment tasks:

- [Network port requirements](#)
- [Network and switch requirements](#)
- [Network cable requirements](#)
- [IP address requirements](#)
- [Network configuration](#)
- [DNS and timekeeping requirements](#)
- [Environmental requirements](#)
- [Protection Domains](#)
- [Witness Node resource requirements for two-node storage clusters](#)

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Network port requirements

You might need to allow the following ports through your data center's edge firewall so that you can manage the system remotely, allow clients outside of your data center to connect to resources, and ensure that internal services can function properly. Some of these ports, URLs, or IP addresses might not be required, depending on how you use the system.

All ports are TCP unless stated otherwise, and all TCP ports must support three-way handshake communication between the NetApp Support Server, management node, and nodes running Element software. For example, the host on a management node source communicates with the host on a storage cluster MVIP destination through TCP port 443, and the destination host communicates back to the source host through any port.

The following abbreviations are used in the table:

- MIP: Management IP address, a per-node address

- SIP: Storage IP address, a per-node address
- MVIP: Management virtual IP address
- SVIP: Storage virtual IP address

Source	Destination	Port	Description
Compute node BMC/IPMI	Management node	111 TCP/UDP	NetApp Hybrid Cloud Control API communication
Compute node BMC/IPMI	Management node	137-138 UDP	NetApp Hybrid Cloud Control API communication
Compute node BMC/IPMI	Management node	445	NetApp Hybrid Cloud Control API communication
Compute node BMC/IPMI	Management node	623 UDP	Remote Management Control Protocol (RMCP) port. Required for NetApp Hybrid Cloud Control compute firmware upgrades.
Compute node BMC/IPMI	Management node	2049 TCP/UDP	NetApp Hybrid Cloud Control API communication
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	Compute node BMC/IPMI	139	NetApp Hybrid Cloud Control API communication
Management node	Storage node MIP	442	UI and API access to storage node and Element software upgrades
Management node	Storage node MVIP	442	UI and API access to storage node and Element software upgrades
Management node	23.32.54.122, 216.240.21.15	443	Element software upgrades

Source	Destination	Port	Description
Management node	Baseboard management controller (BMC)	443	Hardware monitoring and inventory connection (Redfish and IPMI commands)
Management node	Compute node BMC/IPMI	443	NetApp Hybrid Cloud Control HTTPS communication
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
Management node	VMware vCenter	443	NetApp Hybrid Cloud Control HTTPS communication
Management node	Compute node BMC/IPMI	623 UDP	Remote Management Control Protocol (RMCP) port. Required for NetApp Hybrid Cloud Control compute firmware upgrades.
Management node	Storage node BMC/IPMI	623 UDP	RMCP port. This is required to manage IPMI-enabled systems.
Management node	VMware vCenter	5988-5989	NetApp Hybrid Cloud Control HTTPS communication
Management node	Witness Node	9442	Per-node configuration API service
Management node	vCenter Server	9443	vCenter Plug-in registration. The port can be closed after registration is complete.
SNMP server	Storage cluster MVIP	161 UDP	SNMP polling
SNMP server	Storage node MIP	161 UDP	SNMP polling
Storage node BMC/IPMI	Management node	623 UDP	RMCP port. This is required to manage IPMI-enabled systems.
Storage node MIP	DNS servers	53 TCP/UDP	DNS lookup
Storage node MIP	Management node	80	Element software upgrades

Source	Destination	Port	Description
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	Management node	443	Element software upgrades
Storage node MIP	Remote storage cluster MVIP	443	Remote replication cluster pairing communication
Storage node MIP	Remote storage node MIP	443	Remote replication cluster pairing communication
Storage node MIP	S3/Swift endpoint	443	(Optional) HTTPS communication to S3/Swift endpoint for backup and recovery
Storage node MIP	LDAPS server	636 TCP/UDP	LDAPS lookup
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	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	Compute node SIP	442	Compute node API, configuration and validation, and access to software inventory
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

Source	Destination	Port	Description
Storage node SIP	Remote storage node SIP	4000 through 4020	Remote replication node-to-node data transfer
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
System administrator PC	Compute node BMC/IPMI H410 and H600 series	443	HTTPS UI and API access to node remote control
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 BMC/IPMI H410 and H600 series	443	HTTPS UI and API access to node remote control
System administrator PC	Storage node MIP	443	HTTPS storage cluster creation, post-deployment UI access to storage cluster
System administrator PC	Compute node BMC/IPMI H410 and H600 series	623 UDP	RMCP port. This is required to manage IPMI-enabled systems.
System administrator PC	Storage node BMC/IPMI H410 and H600 series	623 UDP	RMCP port. This is required to manage IPMI-enabled systems.
System administrator PC	Witness Node	8080	Witness Node per-node web UI
vCenter Server	Storage cluster MVIP	443	vCenter Plug-in API access
vCenter Server	Management node	8443	(Optional) vCenter Plug-in QoSIOC 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.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Network and switch requirements

The switches you use for NetApp HCI require specific configuration to ensure a successful deployment. See your switch documentation for specific instructions on implementing each of the following requirements for your environment.

A NetApp HCI deployment requires at least three network segments, one for each of the following types of traffic:

- Management
- VMware vMotion
- Storage/Data

Depending on the NetApp H-Series compute and storage node models and the planned cabling configuration, you can physically separate these networks using separate switches or logically separate them using VLANs. For most deployments, however, you need to logically separate these networks (and any other additional virtual machine networks) using VLANs.

Compute and storage nodes need to be able to communicate before, during, and after deployment. If you are implementing separate management networks for storage and compute nodes, ensure that these management networks have network routes between them. These networks must have gateways assigned, and there must be a route between the gateways. Ensure that each new node has a gateway assigned to facilitate communication between nodes and management networks.

NetApp HCI has the following switch requirements:

- All switch ports connected to NetApp HCI nodes must be configured as spanning tree edge ports.
 - On Cisco switches, depending on the switch model, software version and port type, you can do this with one of the following commands:
 - `spanning-tree port type edge`
 - `spanning-tree port type edge trunk`
 - `spanning-tree portfast`
 - `spanning-tree portfast trunk`
 - On Mellanox switches, you can do this with the `spanning-tree port type edge` command.
- NetApp HCI nodes have redundant ports for all network functions except out-of-band management. For the best resiliency, divide these ports across two switches with redundant uplinks to either a traditional hierarchical architecture or a layer 2 spine-and-leaf architecture.
- The switches handling storage, virtual machine, and vMotion traffic must support speeds of at least 10GbE per port (up to 25GbE per port is supported).
- The switches handling management traffic must support speeds of at least 1GbE per port.
- You must configure jumbo frames on the switch ports handling storage and vMotion traffic. Hosts must be able to send 9000 byte packets end-to-end for a successful installation.

- Round-trip network latency between all storage and compute nodes should not exceed 2ms.

All NetApp HCI nodes provide additional out-of-band management capabilities via a dedicated management port. NetApp H300S, H300E, H500S, H500E, H700S, H700E and H410C nodes also allow for IPMI access via Port A. As a best practice, you should ease remote management of NetApp HCI by configuring out-of-band management for all nodes in your environment.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Network cable requirements

You can use the following guidelines to ensure that you have enough of the right type of network cables for the size of your deployment. For RJ45 ports, you must use Cat 5e or Cat 6 rated cables.

- Two-cable compute node configuration: Each compute node must be connected to a 10/25GbE network via two SFP+/SFP28 interfaces (one additional Cat 5e/6 cable is optional for out-of-band management).
- Six-cable compute node configuration: Each compute node must be connected to a 10/25GbE network via four SFP+/SFP28 interfaces and to a 1/10GbE network via two Cat 5e/6 cables (one additional Cat 5e/6 cable is optional for out-of-band management).
- Each storage node must be connected to a 10/25GbE network via two SFP+/SFP28 interfaces and to a 1/10GbE network via two Cat 5e/6 cables (one additional Cat 5e/6 cable is optional for out-of-band management).
- Ensure the network cables you use to connect the NetApp HCI system to your network are long enough to comfortably reach your switches.

For example, a deployment containing four storage nodes and three compute nodes (using the six-cable configuration) requires the following number of network cables:

- (14) Cat 5e/6 cables with RJ45 connectors (plus seven cables for IPMI traffic, if desired)
- (20) Twinax cables with SFP28/SFP+ connectors

This is due to the following reasons:

- Four storage nodes require eight (8) Cat 5e/6 cables and eight (8) Twinax cables.
- Three compute nodes using the six-cable configuration require six (6) Cat 5e/6 cables and twelve (12) Twinax cables.



In a six-cable configuration, two ports are reserved for VMware ESXi and set up and managed by the NetApp Deployment Engine. You cannot access or manage these ESXi-dedicated ports using the Element TUI or the Element web GUI.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

IP address requirements

NetApp HCI has specific IP address requirements that depend on the size of your deployment. Note that by default the initial IP addresses you assign to each node before using the NetApp Deployment Engine to deploy the system are temporary and cannot be reused. You need to set aside a second permanent set of unused IP addresses that you can assign during final deployment.

Number of IP addresses needed per NetApp HCI deployment

The NetApp HCI storage network and management network should each use separate contiguous ranges of IP addresses. Use the following table to determine how many IP addresses you need for your deployment:

System component	Management network IP addresses needed	Storage network IP addresses needed	vMotion network IP addresses needed	Total IP addresses needed per component
Compute node	1	2	1	4
Storage node	1	1		2
Storage cluster	1	1		2
VMware vCenter	1			1
Management node	1	1		2
Witness Node	1	1		2 per Witness Node (two Witness Nodes are deployed for each two-node or three-node storage cluster)

IP addresses reserved by NetApp HCI

NetApp HCI reserves the following IP address ranges for system components. When planning your network, avoid using these IP addresses:

IP address range	Description
10.0.0.0/24	Docker overlay network
10.0.1.0/24	Docker overlay network
10.255.0.0/16	Docker swarm ingress network
169.254.100.1/22	Docker bridge network
169.254.104.0/22	Docker bridge network

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Network configuration

Network configuration

NetApp HCI can utilize multiple different network cabling and VLAN configurations. It is important to plan your network configuration to ensure a successful deployment.

Required network segments

NetApp HCI requires a minimum of three network segments: management, storage, and virtualization traffic (which includes virtual machines and VMware vMotion traffic). You can also separate virtual machine and vMotion traffic. These network segments usually exist as logically separated VLANs in the NetApp HCI network infrastructure.

How compute and storage nodes connect to these networks depends on how you design the network and cable the nodes. The sample network illustrations in this guide assume the following networks:

Network name	VLAN ID
Management	100
Storage	105
vMotion	107
Virtual machines	200, 201

For automatic discovery and configuration of your NetApp HCI nodes in the NetApp Deployment Engine, you must have a network segment that is available as an untagged or native VLAN on all switch ports that are used for the SFP+/SFP28 interfaces on the nodes. This will provide layer 2 communication between all nodes for discovery and deployment. Without a native VLAN, you must configure the SFP+/SFP28 interfaces of all nodes manually with a VLAN and IPv4 address to be discoverable. In the network configuration examples in this document, the management network (VLAN ID 100) is used for this purpose.

The NetApp Deployment Engine enables you to quickly configure networks for compute and storage nodes during the initial deployment. You can place certain built-in management components such as vCenter and the management node on their own network segment. These network segments require routing to allow vCenter and the management node to communicate with storage and compute management networks. In most deployments those components use the same management network (VLAN ID 100 in this example).

 You configure virtual machine networks using vCenter. The default virtual machine network (port group "VM_Network") in NetApp HCI deployments is configured without a VLAN ID. If you plan to use multiple tagged virtual machine networks (VLAN IDs 200 and 201 in the preceding example), ensure you include them in the initial network planning.

Network configuration and cabling options

You can use a two-cable network configuration for the H410C compute nodes, simplifying cable routing. This configuration uses two SFP+/SFP28 interfaces plus an optional (but recommended) RJ45 interface for IPMI communication. These nodes can also use a six-cable configuration with two RJ45 and four SFP28/SFP+ interfaces.

The H410S and H610S storage nodes support a network topology that uses four network ports (ports A through D).

Compute nodes support three network topologies, depending on the hardware platform:

Configuration option	Cabling for H410C nodes	Cabling for H610C nodes	Cabling for H615C nodes
Option A	Two cables using ports D and E	Two cables using ports C and D	Two cables using ports A and B
Option B	Six cables using ports A through F	Not available	Not available
Option C	Similar to option B, but with native VLANs (or "access ports") on the switch for the management, storage, and vMotion networks		

Nodes that do not have the correct number of connected cables cannot participate in the deployment. For example, you cannot deploy a compute node in a six-cable configuration if it only has ports D and E connected.

 You can adjust the NetApp HCI network configuration after deployment to meet infrastructure needs. However, when you expand NetApp HCI resources, remember that new nodes must have the same cable configuration as the existing compute and storage nodes.

 If the NetApp Deployment Engine fails because your network does not support jumbo frames, you can perform one of the following workarounds:

- Use a static IP address and manually set a maximum transmission unit (MTU) of 9000 bytes on the Bond10G network.
- Configure the Dynamic Host Configuration Protocol to advertise an interface MTU of 9000 bytes on the Bond10G network.

Network configuration options

- [Network configuration option A](#)
- [Network configuration option B](#)
- [Network configuration option C](#)

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Network configuration

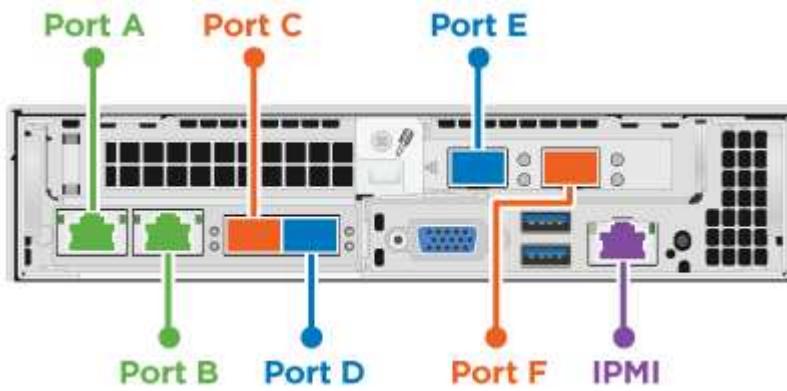
NetApp HCI can utilize multiple different network cabling and VLAN configurations. The first configuration, option A, uses two network cables for each compute node.

Configuration option A: Two cables for compute nodes

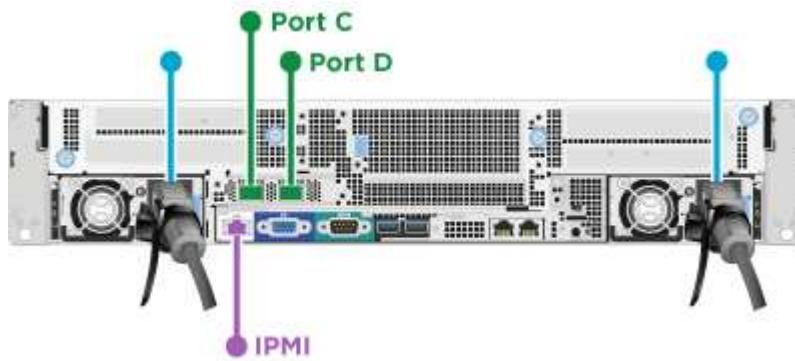
The NetApp H410C, H610C, and H615C compute nodes support using two network cables for connectivity to all NetApp HCI networks. This configuration requires that the storage, vMotion and any virtual machine networks use VLAN tagging. All compute and storage nodes must use the same VLAN ID scheme. This configuration uses vSphere Distributed Switches that require VMware vSphere Enterprise Plus licensing.

NetApp HCI documentation uses letters to refer to network ports on the back panel of H-series nodes.

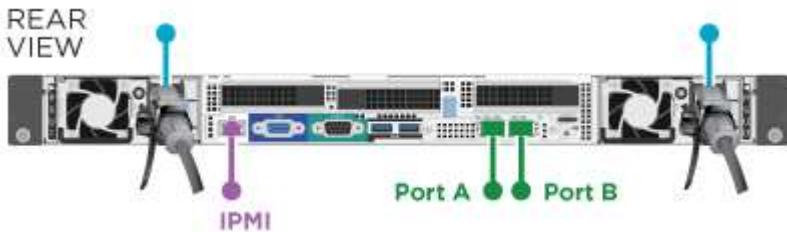
Here are the network ports and locations on the H410C storage node:



Here are the network ports and locations on the H610C compute node:



Here are the network ports and locations on the H615C compute node:



This configuration uses the following network ports on each node:

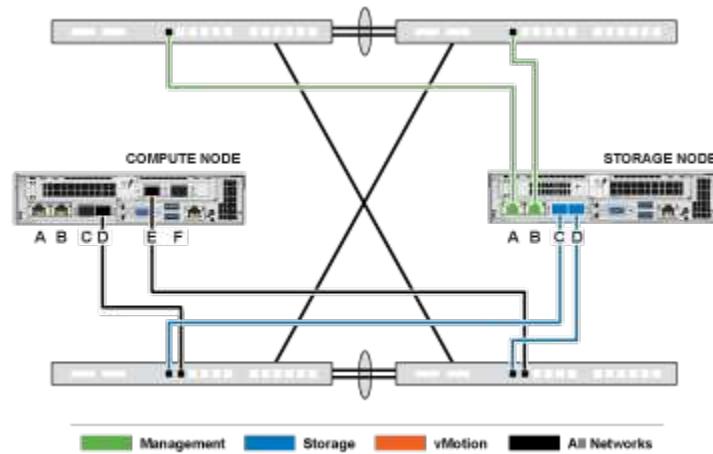
Node	Network ports used
H410C	D and E
H610C	C and D
H615C	A and B

VLAN configuration

As a best practice, you should configure the required network segments on all switch ports that the nodes are using. For example:

Network name	VLAN ID	Switch port configuration
Management	100	Native
Storage	105	Tagged
vMotion	107	Tagged
Virtual machines	200, 201	Tagged

The following illustration shows the recommended cabling configuration for two-cable H410C compute nodes and four-cable H410S storage nodes. All switch ports in this example share the same configuration.



Example switch commands

You can use the following example commands to configure all switch ports used for NetApp HCI nodes. These commands are based on a Cisco configuration, but might require only small changes to apply to Mellanox switches. See your switch documentation for the specific commands you need to implement this configuration. Replace the interface name, description, and VLANs with the values for your environment.

```
interface {interface name, such as EthernetX/Y or GigabitEthernetX/Y/Z}
description {desired description, such as NetApp-HCI-NodeX-PortY}
mtu 9216
switchport mode trunk
switchport trunk native vlan 100
switchport trunk allowed vlan 105,107,200,201
spanning-tree port type edge trunk
```



Some switches might require inclusion of the native VLAN in the allowed VLAN list. See the documentation for your specific switch model and software version.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Network configuration

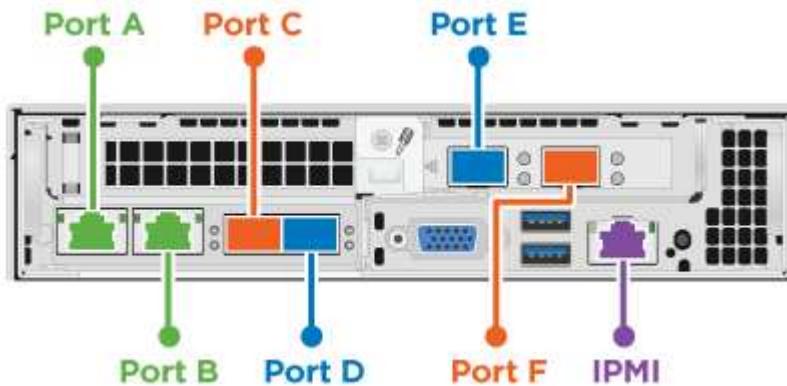
NetApp HCI can utilize multiple different network cabling and VLAN configurations. The first configuration, option B, uses six network cables for each compute node.

Configuration option B: Six cables for compute nodes

As a secondary network configuration option, the H410C compute nodes support using six network cables for connectivity to all NetApp HCI networks. This configuration requires that the storage, vMotion and any virtual machine networks use VLAN tagging. You can use this configuration with vSphere Standard Switches or vSphere Distributed Switches (which require VMware vSphere Enterprise Plus licensing).

NetApp HCI documentation uses letters to refer to network ports on the back panel of H-series nodes.

Here are the network ports and locations on the H410C compute node:

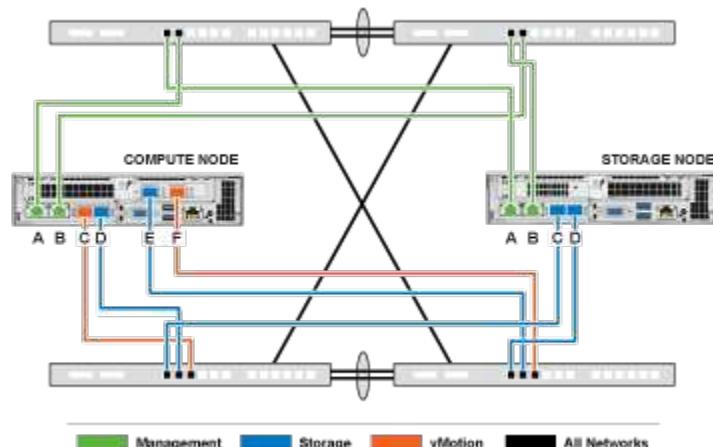


VLAN configuration

When you deploy compute nodes using six cables and storage nodes using four cables, as a best practice, you should configure the required network segments on all switch ports that the nodes are using. For example:

Network name	VLAN ID	Switch port configuration
Management	100	Native
Storage	105	Tagged
vMotion	107	Tagged
Virtual machines	200, 201	Tagged

The following illustration shows the recommended cabling configuration for six-cable compute nodes and four-cable storage nodes. All switch ports in this example share the same configuration.



Example switch commands

You can use the following example commands to configure all switch ports used for NetApp HCI nodes. These commands are based on a Cisco configuration, but might require only small changes to apply to Mellanox switches. See your switch documentation for the specific commands you need to implement this configuration. Replace the interface name, description, and VLANs with the values for your environment.

```
interface {interface name, such as EthernetX/Y or GigabitEthernetX/Y/Z}
description {desired description, such as NetApp-HCI-NodeX-PortY}
mtu 9216
switchport mode trunk
switchport trunk native vlan 100
switchport trunk allowed vlan 105,107,200,201
spanning-tree port type edge trunk
```



Some switches might require inclusion of the native VLAN in the allowed VLAN list. See the documentation for your specific switch model and software version.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Network configuration

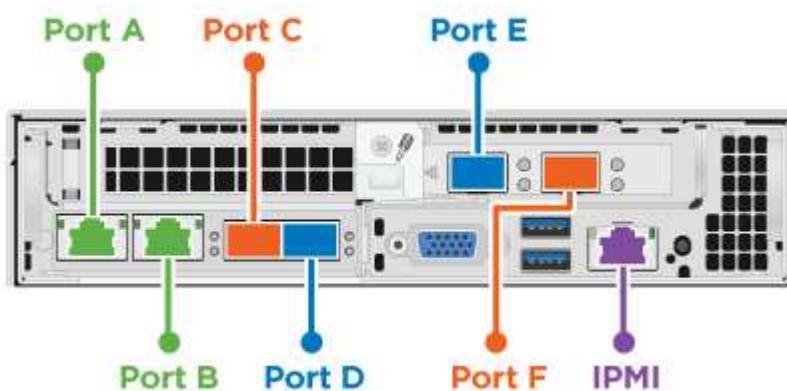
NetApp HCI can utilize multiple different network cabling and VLAN configurations. The third configuration, option C, uses six network cables for each compute node with native VLANs.

Configuration option C: Six cables for compute nodes with native VLANs

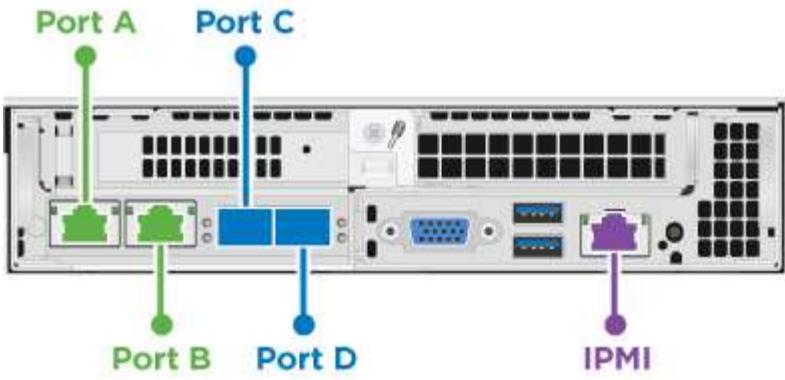
You can deploy NetApp HCI without using tagged VLANs for storage and virtualization traffic, and instead rely on the switch configuration to separate the network segments. You can use this configuration with vSphere Standard Switches or vSphere Distributed Switches (which require VMware vSphere Enterprise Plus licensing).

NetApp HCI documentation uses letters to refer to network ports on the back panel of H-series nodes.

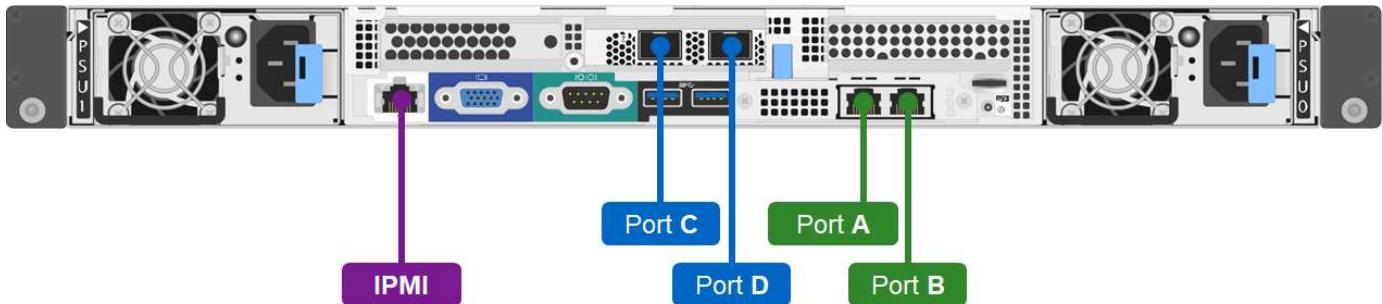
Here are the network ports and locations on the H410C storage node:



Here are the network ports and locations on the H410S storage node:



Here are the network ports and locations on the H610S storage node:



VLAN configuration for H410C, H410S, and H610S nodes

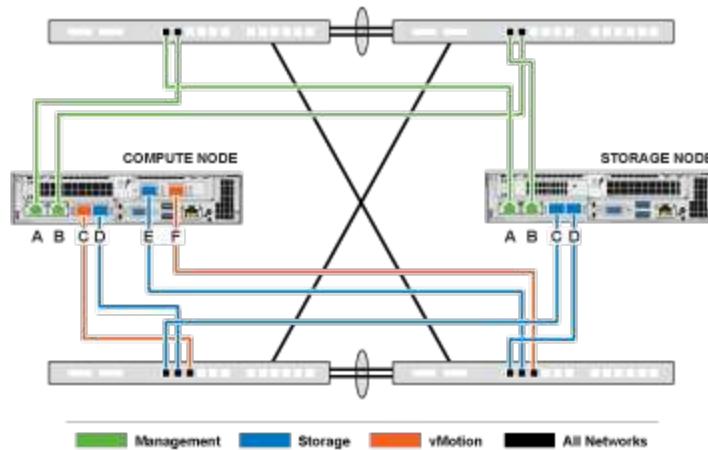
This topology option uses the following VLAN configuration on H410C, H410S, and H610S nodes:

Node ports used	Network name	VLAN ID	Connected switch port configuration
Ports A and B on compute and storage nodes	Management	100	Native
Ports D and E on compute nodes	Storage	105	Native
Ports C and D on storage nodes	Storage	105	Native
Ports C and F on compute nodes	vMotion	107	Native
Ports C and F on compute nodes	Virtual machines	200, 201	Tagged



Be careful configuring the switch ports when deploying this configuration. Configuration errors in this network topology can result in deployment problems that are difficult to diagnose.

The following illustration shows the network configuration overview for this topology option. In the example, individual switch ports are configured with the appropriate network segment as the native network.



Example switch commands

You can use the following example switch commands to configure switch ports used for the NetApp HCI nodes. These commands are based on a Cisco configuration, but might require only minimal changes to apply to Mellanox switches. See your switch documentation for the specific commands you need to implement this configuration.

You can use the following example commands to configure the switch ports used for the management network. Replace the interface name, description, and VLANs with the values for your configuration.

```
interface {interface name, such as EthernetX/Y or GigabitEthernetX/Y/Z}
description {desired description, such as NetApp-HCI-NodeX-PortA|B}
switchport access vlan 100
spanning-tree port type edge
```

You can use the following example commands to configure the switch ports used for the storage network. Replace the interface name, description, and VLANs with the values for your configuration.

```
interface {interface name, such as EthernetX/Y or GigabitEthernetX/Y/Z}
description {desired description, such as NetApp-HCI-NodeX-PortC|D}
mtu 9216
switchport access vlan 105
spanning-tree port type edge
```

You can use the following example commands to configure the switch ports used for the vMotion and virtual machines network. Replace the interface name, description, and VLANs with the values for your configuration.

```
interface {interface name, such as EthernetX/Y or GigabitEthernetX/Y/Z}
description {desired description, such as NetApp-HCI-NodeX-PortC|F}
mtu 9216
switchport mode trunk
switchport trunk native vlan 107
switchport trunk allowed vlan 200,201
spanning-tree port type edge trunk
```



Some switches might require inclusion of the native VLAN in the allowed VLAN list. See the documentation for your specific switch model and software version.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

DNS and timekeeping requirements

Before deployment, you need to prepare Domain Name System (DNS) records for your NetApp HCI system and gather NTP server information. NetApp HCI requires a DNS server with the correct DNS entries and an NTP server for a successful deployment.

Make the following DNS and timeserver preparations before deploying NetApp HCI:

- Create any needed DNS entries for hosts (such as individual compute or storage nodes) and document how the host entries map to the respective IP addresses. During deployment, you will need to assign a prefix to your storage cluster that will be applied to each host; to avoid confusion, keep your DNS naming plans in mind when choosing a prefix.
- If you are deploying NetApp HCI with a new VMware vSphere installation using a fully qualified domain name, you must create one Pointer (PTR) record and one Address (A) record for vCenter Server on any DNS servers in use before deployment.
- If you are deploying NetApp HCI with a new vSphere installation using only IP addresses, you do not need to create new DNS records for vCenter.
- NetApp HCI requires a valid NTP server for timekeeping. You can use a publicly available time server if you do not have one in your environment.
- Ensure that all storage and compute node clocks are in sync with each other, and that the clocks of devices you use to log in to NetApp HCI are in sync with the NetApp HCI nodes.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Environmental requirements

Ensure that the power for the rack used to install NetApp HCI is supplied by AC power outlets, and that your datacenter provides adequate cooling for the size of your NetApp HCI installation.

For detailed capabilities of each component of NetApp HCI, see the NetApp HCI [datasheet](#).



The H410C compute node operates only on high-line voltage (200-240 VAC). You must ensure that the power requirements are met when you add H410C nodes to an existing NetApp HCI installation.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Protection domains

NetApp Element software supports protection domains functionality, which optimizes data layout on storage nodes for the best data availability. To use this feature, you should split storage capacity evenly across three or more NetApp H-series chassis for optimal storage reliability. In this scenario, the storage cluster automatically enables protection domains.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Witness Node resource requirements for two-node storage clusters

NetApp HCI supports a minimum installation size of two storage nodes and two compute nodes. When you install NetApp HCI using a two or three-node storage cluster, you need to be aware of NetApp HCI Witness Nodes and their virtual machine (VM) resource requirements.

When a storage cluster uses two or three nodes, it also deploys a pair of Witness Nodes alongside each storage cluster. Witness Nodes have the following VM resource requirements:

Resource	Requirement
vCPU	4
Memory	12GB
Disk size	67GB

NetApp HCI supports only certain storage node models in two-node or three-node storage clusters. For more information, see the Release Notes for your NetApp HCI version.

Best practice: Configure the Witness Node VMs to use the compute node's local datastore (default set by NDE), do not configure them on shared storage, such as SolidFire storage volumes. To prevent the VMs migrating automatically, set the Witness Node VM's Distributed Resource Scheduler (DRS) automation level to **Disabled**. This prevents both Witness Nodes running on the same compute node and creating a non-high availability (HA) pair configuration.



When the NetApp HCI installation process installs Witness Nodes, a VM template is stored in VMware vCenter that you can use to redeploy a Witness Node in case it is accidentally removed, lost, or corrupted. You can also use the template to redeploy a Witness Node if you need to replace a failed compute node that was hosting the Witness Node. For instructions, see the **Redeploy Witness Nodes for two and three-node storage clusters** section [here](#).

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Get started with NetApp HCI

NetApp HCI installation and deployment overview

Use these instructions to install and deploy NetApp HCI. These instructions include links to more details.

Here is an overview of the process:

- Prepare for installation
- Validate network readiness with NetApp Active IQ Config Advisor
- Work with your NetApp team
- Install NetApp HCI hardware
- Complete optional tasks after installing hardware
- Deploy NetApp HCI using the NetApp Deployment Engine (NDE)
- Manage NetApp HCI using the vCenter Plug-in
- Monitor or upgrade NetApp HCI with the Hybrid Cloud Control

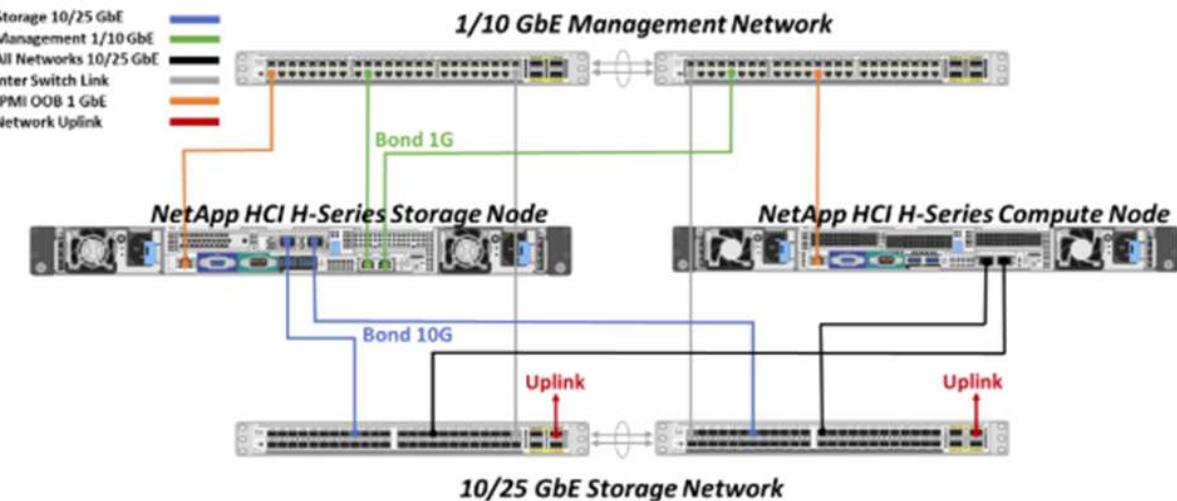
Prepare for installation

Before you begin the installation, complete the *NetApp HCI Installation Discovery Workbook* pre-flight checklist sent to you prior to receiving the hardware.

Prepare the network and installation sites

Here is a simplified NetApp HCI network topology installation:

NetApp HCI Simplified Network Topology Installation



This is the simplified network topology for a single storage node and single compute node. The minimum cluster for NetApp HCI is two storage and two compute nodes.



Your network topology might differ from what is shown here. This is an example only.

This setup uses two network cables on the compute nodes for connectivity to all NetApp HCI networks.

Read these resources:

- Use the *NetApp HCI Installation Discovery Workbook* to configure your network before the installation.
- For details and other supported configurations, see [TR-4820: NetApp HCI Networking Quick Planning Guide](#) and the [NetApp HCI Installation and Setup Instructions](#).
- For information about NetApp HCI configurations smaller than four storage nodes, see [TR-4823: NetApp HCI 2-Node Storage Cluster](#).
- For details about configuring Link Aggregation Control Protocol (LACP) on the switch ports used for each of the storage nodes, see [Configure LACP for optimal storage performance](#).

This setup consolidates all traffic onto two physical, redundant ports, reducing the cabling and streamlining network configuration. This configuration requires that the storage, vMotion and any virtual machine network segments use VLAN tagging. The management network segment can use native or tagged VLAN; however, native VLAN is the preferred mode so that NetApp Deployment Engine (NDE) can assign network resources in an automated manner (Zero Conf).

This mode requires vSphere Distributed Switches (vDS), which require VMware vSphere Enterprise Plus licensing.

Networking requirements before you begin

Here are highlights of prerequisites.

For prerequisites details, see [Requirements for NetApp HCI deployment overview](#).

- Bond1G is a logical interface that combines 1GbE network ports on storage nodes and a management interface on compute nodes. This network is used for NDE API traffic. All nodes must be able to communicate over the management interface in the same L2 network.
- Bond10G is a logical interface that combines 10/25GbE ports and are used by NDE for beaconing and inventory. All nodes must be able to communicate over the Bond10G interface with non-fragmented jumbo frames.
- NDE requires at a minimum one manually assigned IP address on the Bond1G interface on one storage node. NDE will be run from this node.
- All nodes will have temporary IP addresses assigned by NDE discovery, which is accomplished by Automatic Private IP Addressing (APIPA).



During the NDE process, all nodes will then be assigned permanent IP addresses and any APIPA assigned temporary IPs will be released.

- NDE requires separate networks for management, iSCSI and vMotion that are preconfigured on the switch network.

Validate network readiness with NetApp Active IQ Config Advisor

To ensure network readiness for NetApp HCI, install the NetApp Active IQ Config Advisor 5.8.1 or later. This network validation tool is located with other [NetApp Support Tools](#). Use this tool to validate connectivity, VLAN IDs, IP address requirements, switch connectivity and more.

For details, see [Validate your environment with Active IQ Config Advisor](#)

Work with your NetApp team

Your NetApp team uses the NetApp Active IQ Config Advisor report and the *Discovery Workbook* to validate that your network environment is ready.

Install NetApp HCI hardware

NetApp HCI can be installed in different configurations:

- H410C compute nodes: Two-cable configuration or six-cable configuration
- H610C compute node: Two-cable configuration
- H615C compute node: Two-cable configuration
- H410S storage node
- H610S storage node



For precautions and details, see [Install H-series hardware](#).

Steps

1. Install the rails and the chassis.
2. Install nodes in the chassis and install drives for storage nodes. (Applies only if you are installing H410C and H410S in a NetApp H-series chassis.)
3. Install the switches.
4. Cable the compute node.
5. Cable the storage node.
6. Connect the power cords.
7. Power on the NetApp HCI nodes.

Complete optional tasks after installing hardware

After installing the NetApp HCI hardware, you should perform some optional, yet recommended tasks.

Manage storage capacity across all chassis

Ensure that storage capacity is split evenly across all chassis containing storage nodes.

Configure IPMI for each node

After you have racked, cabled, and powered on your NetApp HCI hardware, you can configure Intelligent Platform Management Interface (IPMI) access for each node. Assign each IPMI port an IP address and change the default administrator IPMI password as soon as you have remote IPMI access to the node.

See [Configure IPMI](#).

Deploy NetApp HCI using the NetApp Deployment Engine (NDE)

The NDE UI is the software wizard interface used to install NetApp HCI.

Launch the NDE UI

NetApp HCI uses a storage node management network IPv4 address for initial access to the NDE. As a best practice, connect from the first storage node.

Prerequisites

- You already assigned the initial storage node management network IP address manually or by using DHCP.
- You must have physical access to the NetApp HCI installation.

Steps

1. If you do not know the initial storage node management network IP, use the Terminal User Interface (TUI), which is accessed via keyboard and monitor on the storage node or [use a USB stick](#).

For details, see [Accessing the NetApp Deployment Engine](#).

2. If you do know the IP address, from a web browser, connect to the Bond1G address of the primary node via HTTP, not HTTPS.

Example: http://<IP_address>:442/nde/

Deploy NetApp HCI with the NDE UI

1. In the NDE, accept the prerequisites, check to use Active IQ, and accept license agreements.
2. Optionally, enable Data Fabric File Services by ONTAP Select and accept the ONTAP Select license.
3. Configure a new vCenter deployment. Select **Configure Using a Fully Qualified Domain Name** and enter both the vCenter Server Domain Name and DNS Server IP address.



It is strongly recommended to use the FQDN approach for vCenter installation.

4. Review that the inventory assessment of all nodes completed successfully.

The storage node that is running the NDE is already checked.

5. Select all nodes and select **Continue**.
6. Configure network settings. Refer to the *NetApp HCI Installation Discovery Workbook* for the values to use.
7. Select the blue box to launch the easy form.

Network Settings

Provide the network settings that will be used for your installation.

Live network validation is: **On** 

Infrastructure Services

DNS Server IP Address 1

DNS Server IP Address 2 (Optional)

NTP Server Address 1 

NTP Server Address 2 (Optional)

To save time, launch the easy form to enter fewer network settings. 

vCenter Networking

VLAN ID	Subnet 	Default Gateway	FQDN	IP Address
Untagged Network	100.0.0.0/16			

8. On the Network Settings Easy Form:

- Type the Naming Prefix. (Refer to the System Details of the *NetApp HCI Installation Discovery Workbook*.)
- Select **No** for Will you assign VLAN IDs? (You assign them later in the main Network Settings page.)
- Type the subnet CIDR, default gateway, and starting IP address for the management, vMotion, and iSCSI networks according to your workbook. (Refer to the IP Assignment Method section of the *NetApp HCI Installation Discovery Workbook* for these values.)
- Select **Apply to Network Settings**.

9. Join an [existing vCenter](#) (optional).

10. Record node serial numbers in the *NetApp HCI Installation Discovery Workbook*.

11. Specify a VLAN ID for the vMotion Network and any network that requires VLAN tagging. See the *NetApp HCI Installation Discovery Workbook*.

12. Download your configuration as a .CSV file.

13. Select **Start Deployment**.

14. Copy and save the URL that appears.



It can take about 45 minutes to complete the deployment.

Verify the installation using the vSphere Web Client

1. Launch the vSphere Web Client and log in using the credentials specified during NDE use.

You must append @vsphere.local to the user name.

2. Verify that no alarms are present.

3. Verify that the vCenter, mNode, and ONTAP Select (optional) appliances are running without warning

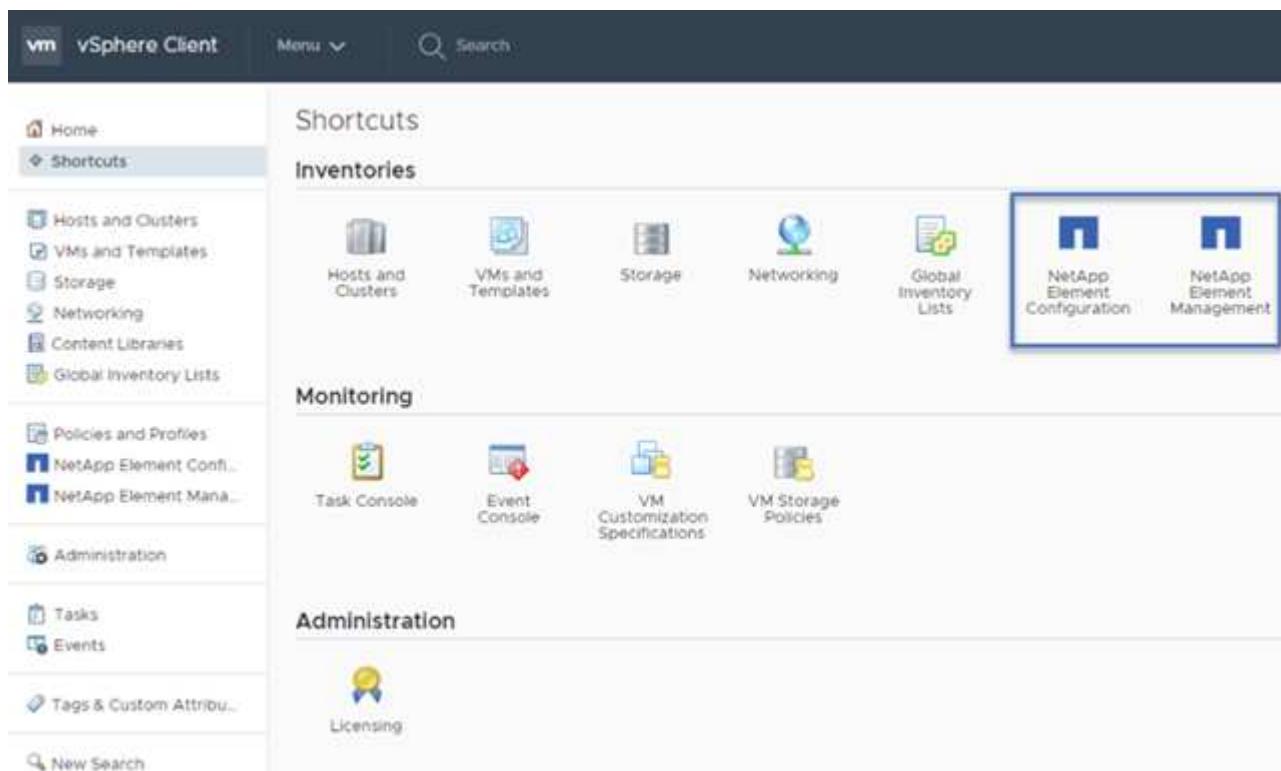
icons.

4. Observe that the two default datastores (NetApp-HCI-Datastore_01 & 02) are created.
5. Select each datastore and ensure that all compute nodes are listed in the Hosts tab.
6. Validate vMotion and Datastore-02.
 - a. Migrate the vCenter Server to NetApp-HCI-Datastore-02 (storage only vMotion).
 - b. Migrate the vCenter Server to each of the compute nodes (compute only vMotion).
7. Go to the NetApp Element Plug-in for vCenter Server and ensure that the cluster is visible.
8. Ensure no alerts appear on the Dashboard.

Manage NetApp HCI using the vCenter Plug-in

After you install NetApp HCI, you can configure clusters, volumes, datastores, logs, access groups, initiators, and Quality of Service (QoS) policies by using the NetApp Element Plug-in for vCenter Server.

For details, see [NetApp Element Plug-in for vCenter Server documentation](#).



Monitor or upgrade NetApp HCI with the Hybrid Cloud Control

You can optionally use the NetApp HCI Hybrid Cloud Control to monitor, upgrade, or expand your system.

You log in to NetApp Hybrid Cloud Control by browsing to the IP address of the management node.

Using the Hybrid Cloud Control, you can do the following:

- Monitor your NetApp HCI installation
- Upgrade your NetApp HCI system
- Expand your NetApp HCI storage or compute resources

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.

The NetApp Hybrid Cloud Control interface appears.

Find more information

- [NetApp HCI Installation and Setup Instructions](#)
- [TR-4820: NetApp HCI Networking Quick Planning Guide](#)
- [NetApp Element Plug-in for vCenter Server Guide](#)
- [NetApp Configuration Advisor 5.8.1 or later network validation tool](#)
- [NetApp SolidFire Active IQ Documentation](#)

Install H-series hardware

Before you get started with using NetApp HCI, you should install the storage and compute nodes correctly.



See the [poster](#) for a visual representation of the instructions.

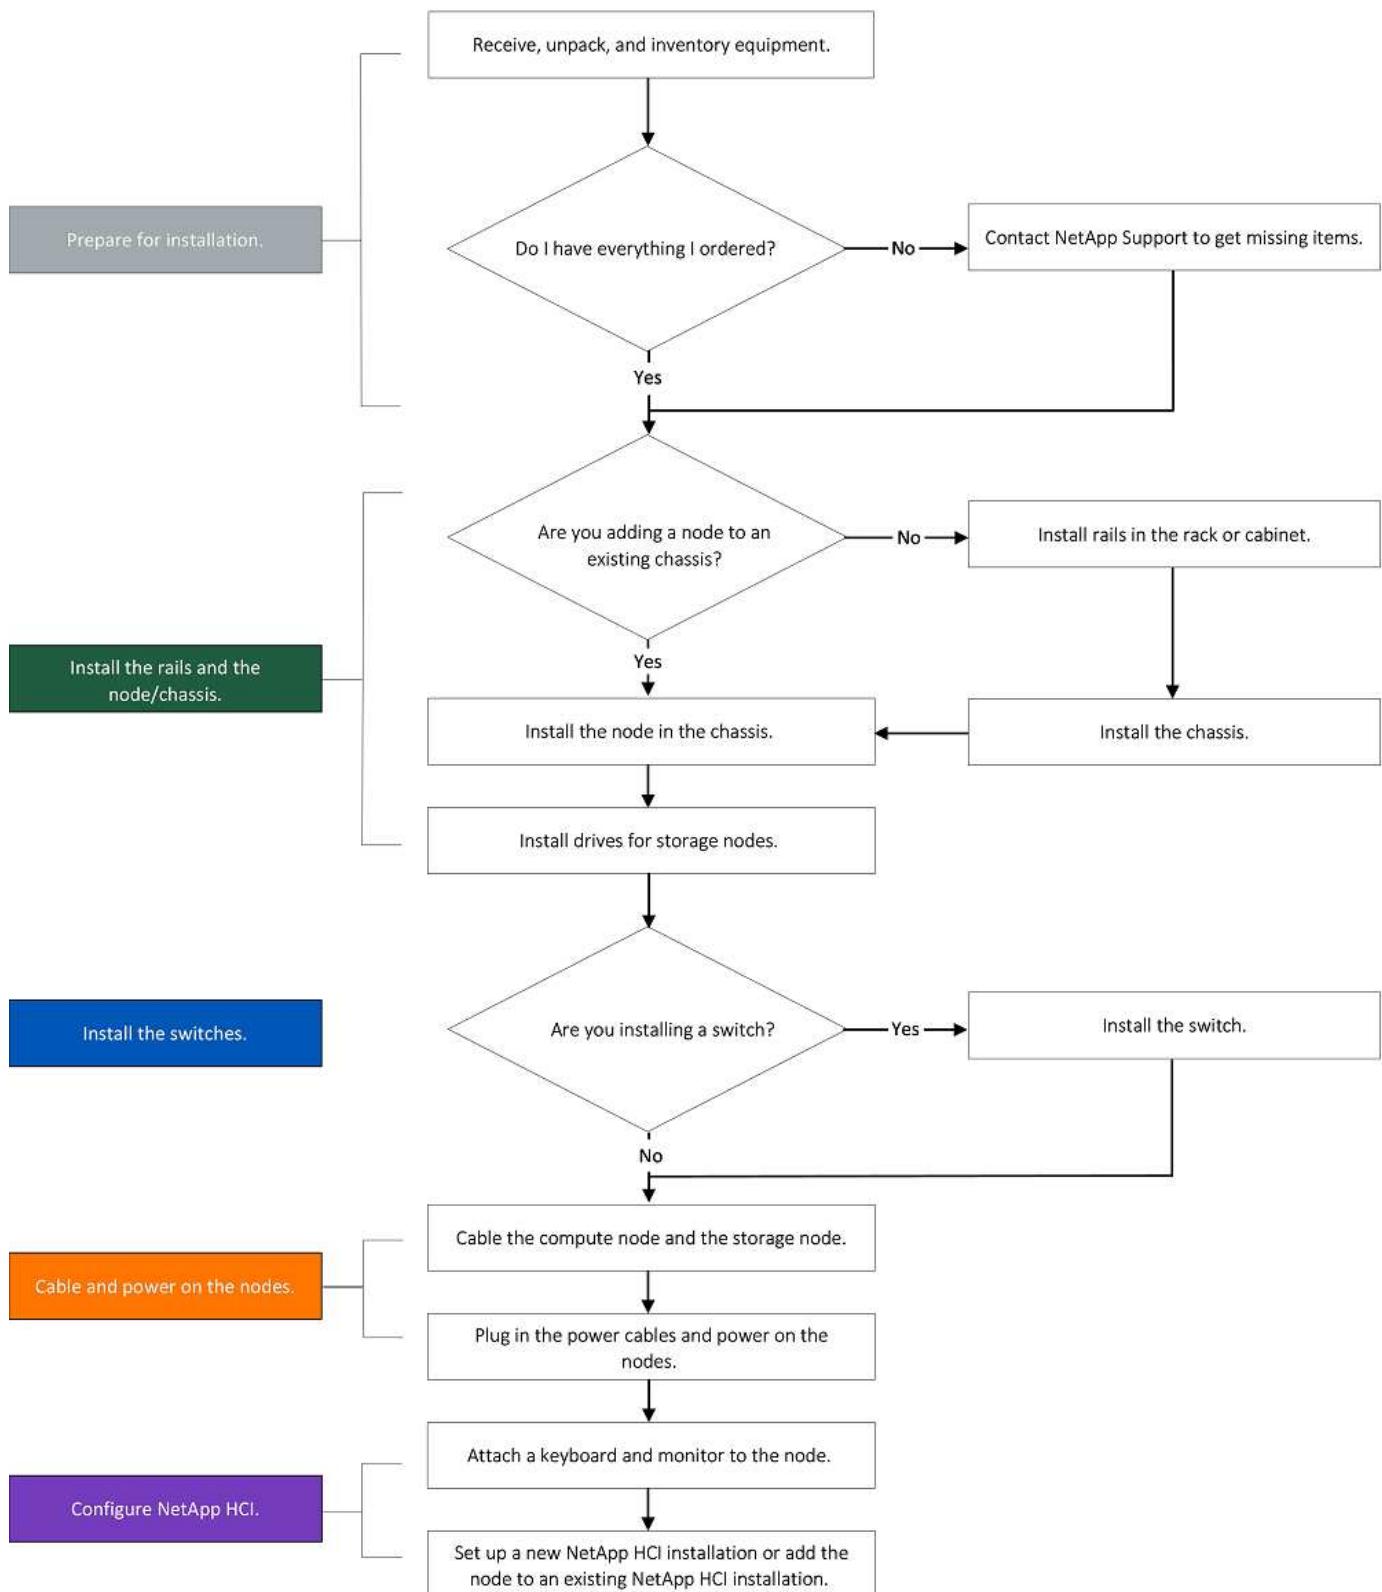
- [Workflow diagrams](#)
- [Prepare for installation](#)
- [Install the rails](#)
- [Install the node/chassis](#)
- [Install the switches](#)
- [Cable the nodes](#)
- [Power on the nodes](#)
- [Configure NetApp HCI](#)
- [Perform post-configuration tasks](#)

Workflow diagrams

The workflow diagrams here provide a high-level overview of the installation steps. The steps vary slightly depending on the H-series model.

- [H410C and H410S](#)
- [H610C and H615C](#)
- [\[H610S\]](#)

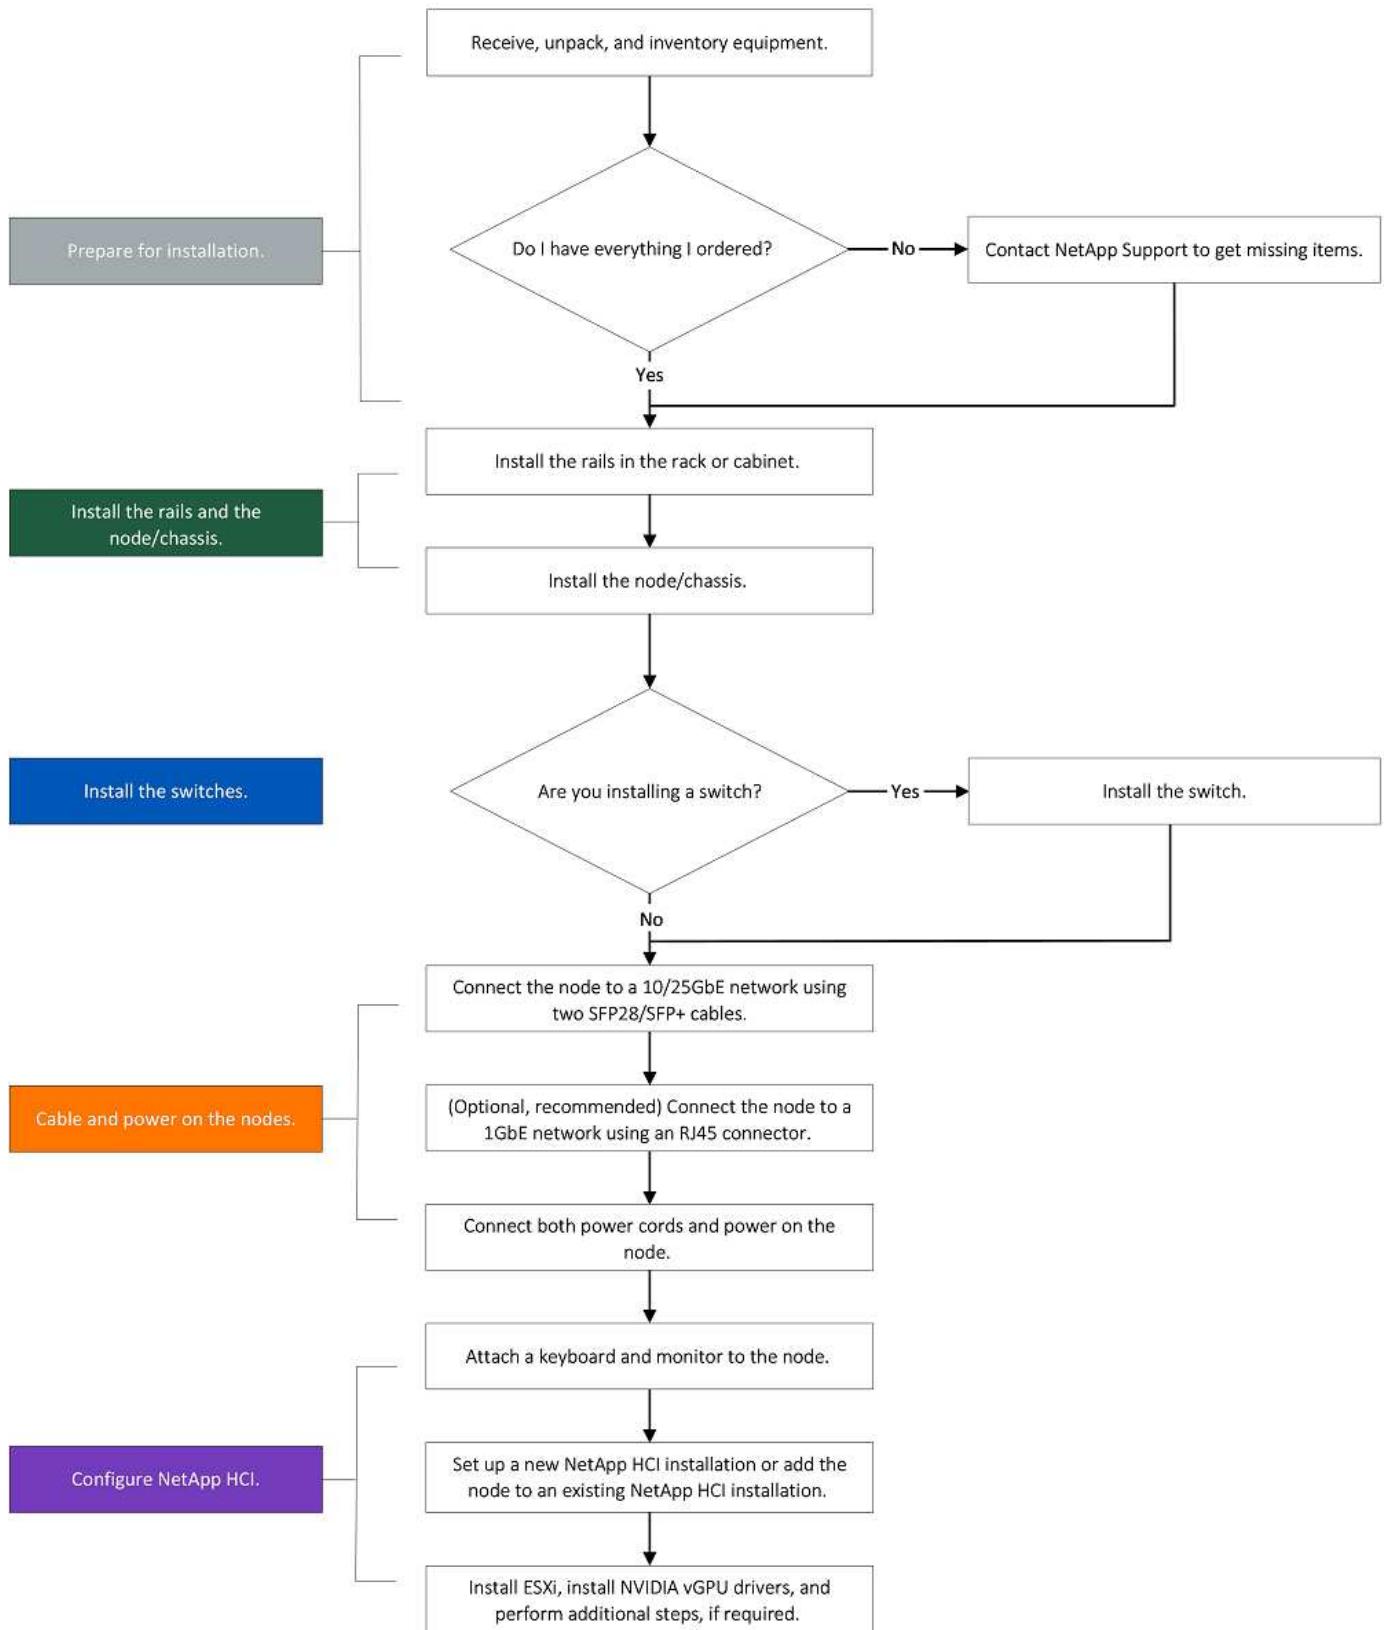
H410C and H410S



H610C and H615C

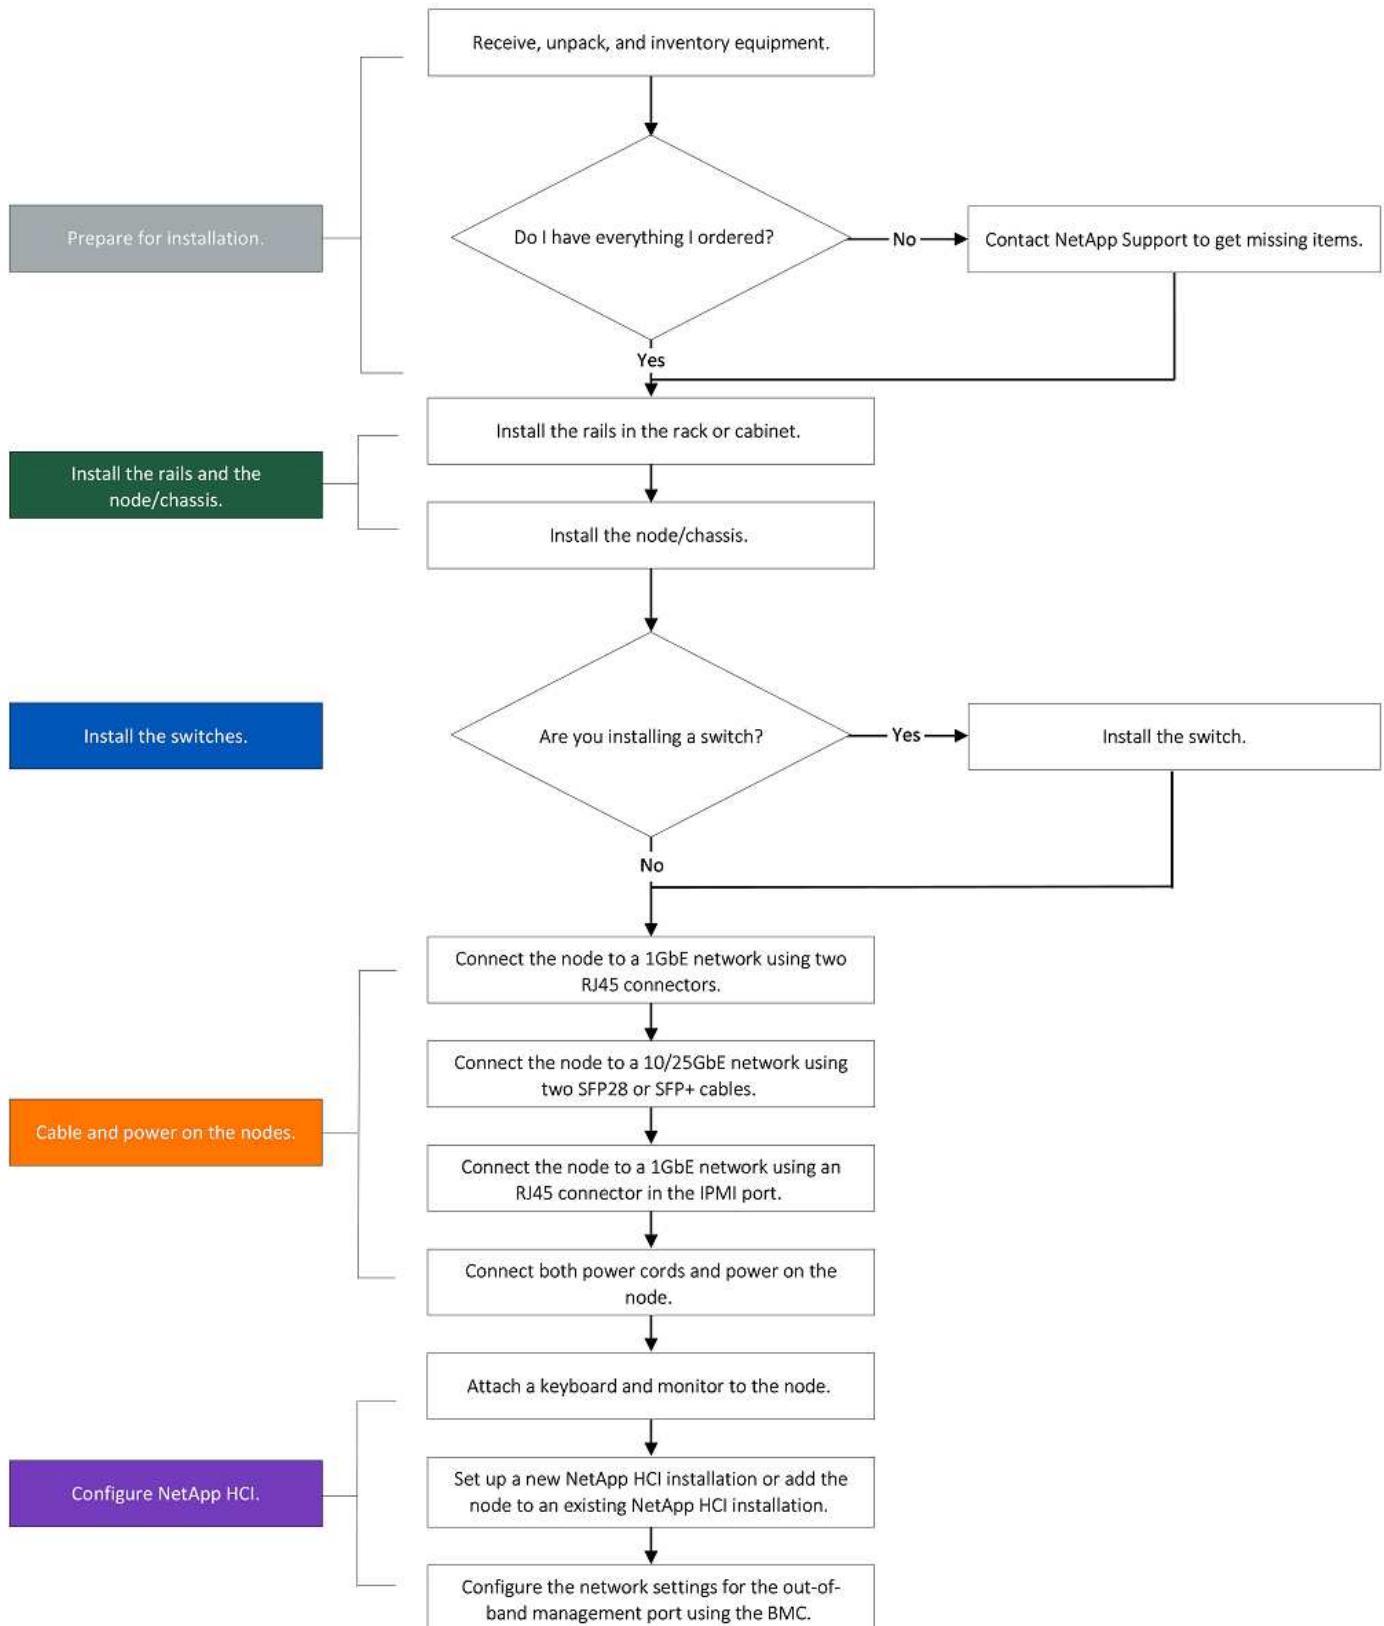


The terms "node" and "chassis" are used interchangeably in the case of H610C and H615C, because node and chassis are not separate components unlike in the case of a 2U, four-node chassis.





The terms "node" and "chassis" are used interchangeably in the case of H610C and H615C, because node and chassis are not separate components unlike in the case of a 2U, four-node chassis.



Prepare for installation

In preparation for installation, inventory the hardware that was shipped to you, and contact NetApp Support if any of the items are missing.

Ensure that you have the following items at your installation location:

- Rack space for the system.

Node type	Rack space
H410C and H410S nodes	Two rack unit (2U)
H610C node	2U
H615C and H610S nodes	One rack unit (1U)

- SFP28/SFP+ direct-attach cables or transceivers
- CAT5e or higher cables with RJ45 connector
- A keyboard, video, mouse (KVM) switch to configure your system
- USB stick (optional)

 The hardware that is shipped to you depends on what you order. A new 2U, four-node order includes the chassis, bezel, slide rail kit, drives for storage nodes, storage and compute nodes, and power cables (two per chassis). If you order H610S storage nodes, the drives will come installed in the chassis.

 While installing the hardware, ensure that you remove all packing material and wrapping from the unit. This will prevent the nodes from overheating and shutting down.

Install the rails

The hardware order that was shipped to you includes a set of slide rails. You will need a screwdriver to complete the rail installation. The installation steps vary slightly for each node model.

 Install hardware from the bottom of the rack up to the top to prevent the equipment from toppling over. If your rack includes stabilizing devices, install them before you install the hardware.

- [H410C and H410S](#)
- [\[H610C\]](#)
- [H610S and H615C](#)

H410C and H410S

H410C and H410S nodes are installed in 2U, four-node H-Series chassis, which is shipped with two sets of adapters. If you want to install the chassis in a rack with round holes, use the adapters appropriate for a rack with round holes. The rails for H410C and H410S nodes fit a rack between 29 inches and 33.5 inches in depth. When the rail is fully contracted, it is 28 inches long, and the front and rear sections of the rail are held together by only one screw.



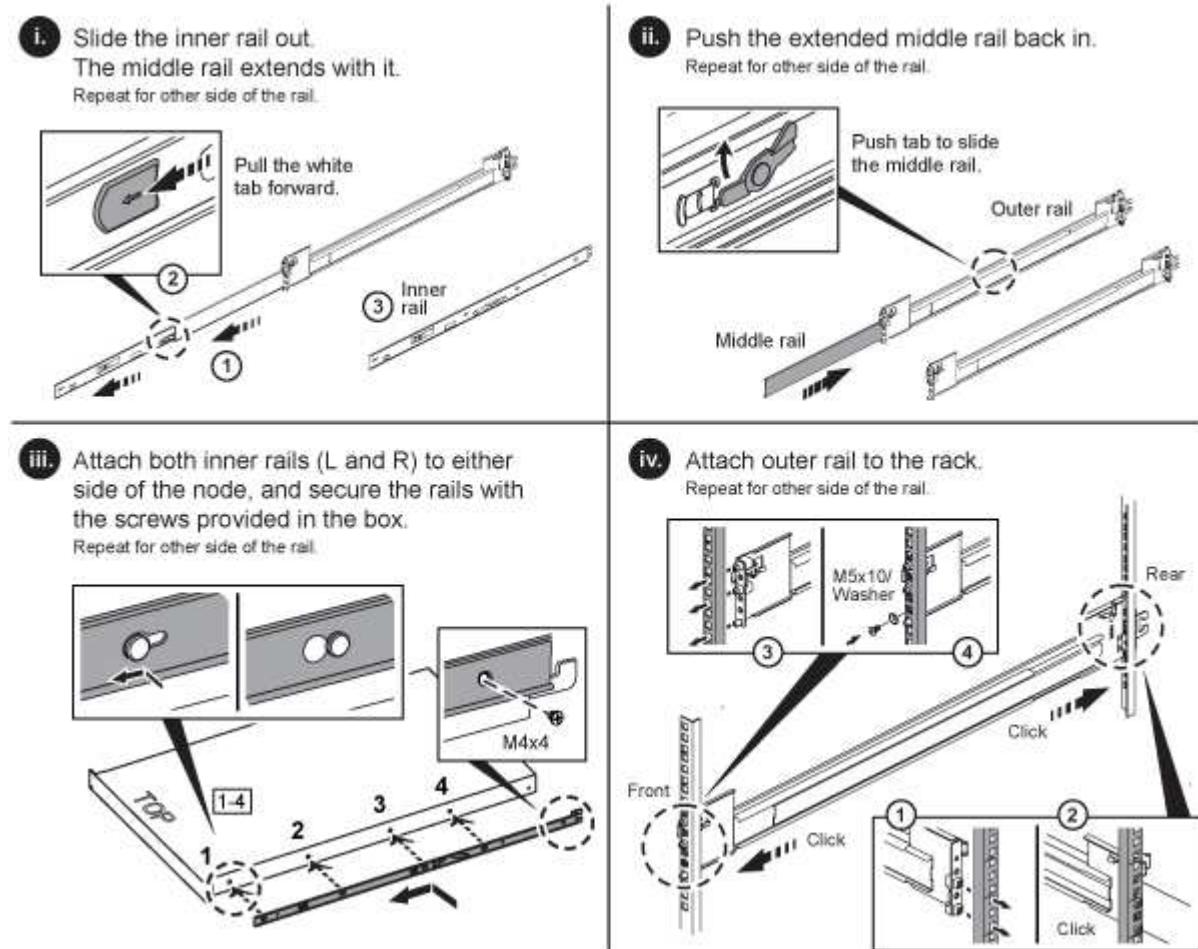
If you install the chassis onto a fully contracted rail, the front and rear sections of the rail might separate.

Steps

1. Align the front of the rail with the holes on the front post of the rack.
 2. Push the hooks on the front of the rail into the holes on the front post of the rack and then down, until the spring-loaded pegs snap into the rack holes.
 3. Attach the rail to the rack with screws. Here is an illustration of the left rail being attached to the front of the rack:
 4. Extend the rear section of the rail to the rear post of the rack.
 5. Align the hooks on the rear of the rail with the appropriate holes on the rear post ensuring that the front and the back of the rail are on the same level.
 6. Mount the rear of the rail onto the rack, and secure the rail with screws.
 7. Perform all the above steps for the other side of the rack.

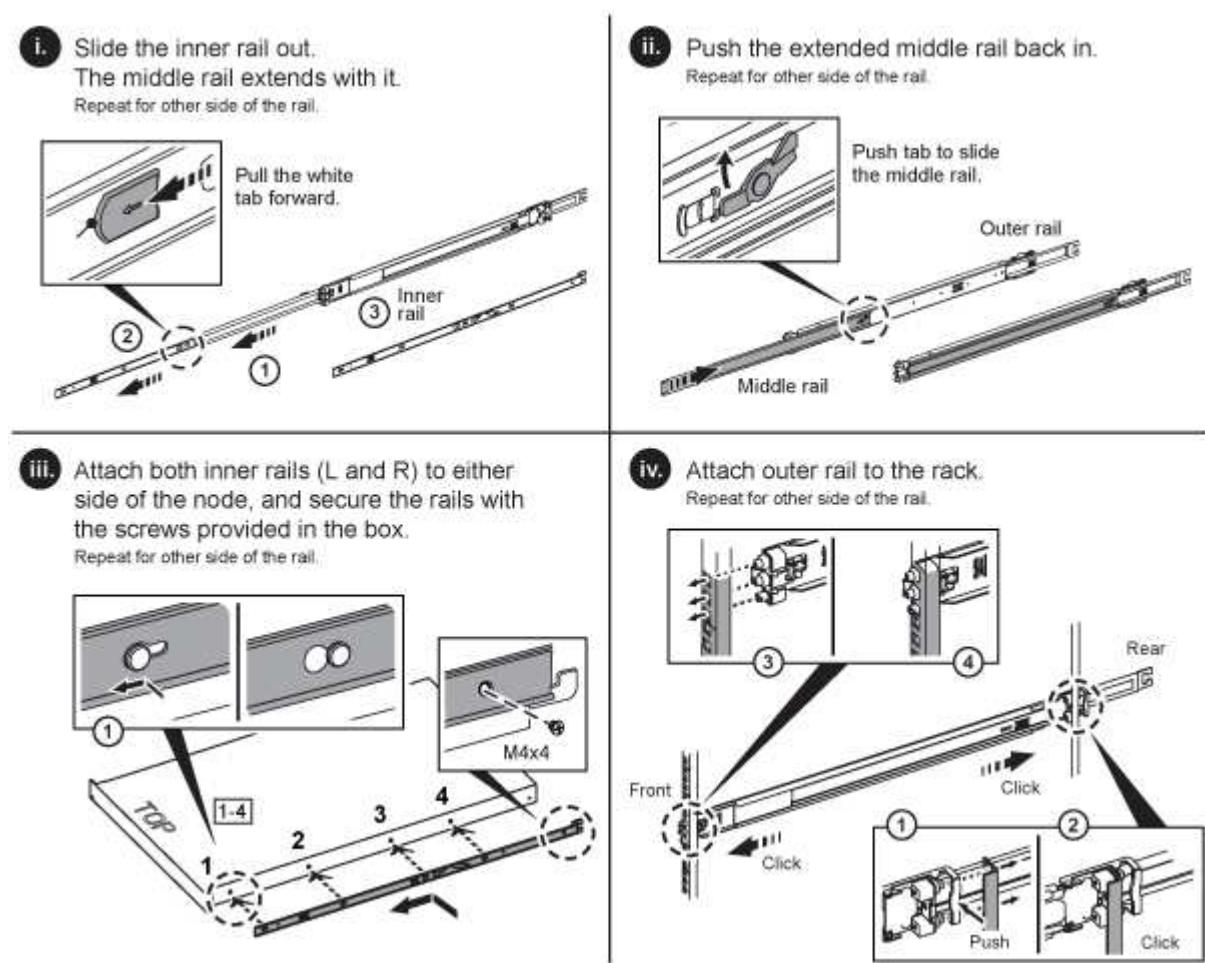
H610C

Here is an illustration for installing rails for an H610C compute node:



H610S and H615C

Here is an illustration for installing rails for an H610S storage node or an H615C compute node:



There are left and right rails on the H610S and H615C. Position the screw hole towards the bottom so that the H610S/H615C thumbscrew can secure the chassis to the rail.

Install the node/chassis

You install the H410C compute node and H410S storage node in a 2U, four-node chassis. For H610C, H615C, and H610S, install the chassis/node directly onto the rails in the rack.



Starting with NetApp HCI 1.8, you can set up a storage cluster with two or three storage nodes.



Remove all the packing material and wrapping from the unit. This prevents the nodes from overheating and shutting down.

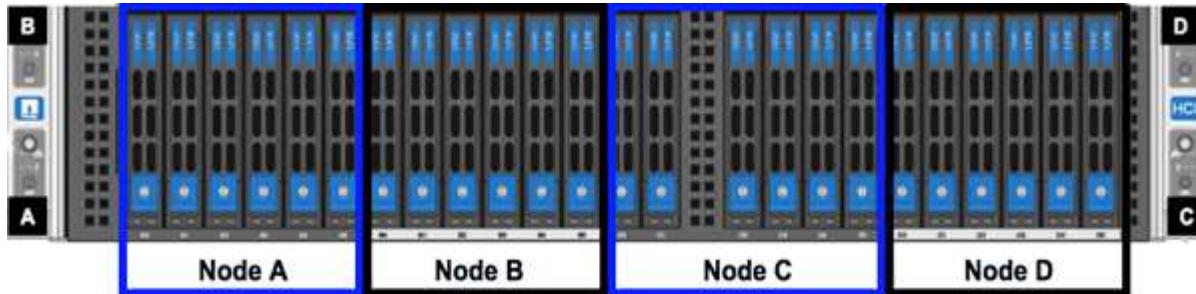
- [H410C and H410S nodes](#)
- [H610C node/chassis](#)
- [H610S and H615C node/chassis](#)

H410C and H410S nodes

Steps

1. Install the H410C and H410S nodes in the chassis. Here is a rear-view example of a chassis with four nodes installed:

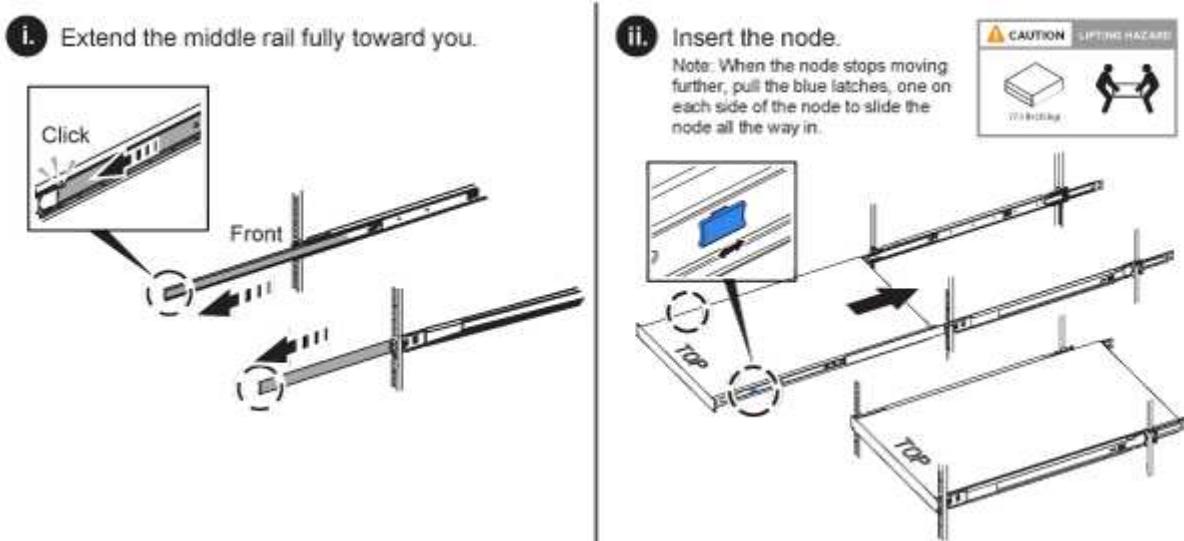
2. Install drives for H410S storage nodes.



H610C node/chassis

In the case of H610C, the terms "node" and "chassis" are used interchangeably because node and chassis are not separate components, unlike in the case of the 2U, four-node chassis.

Here is an illustration for installing the node/chassis in the rack:

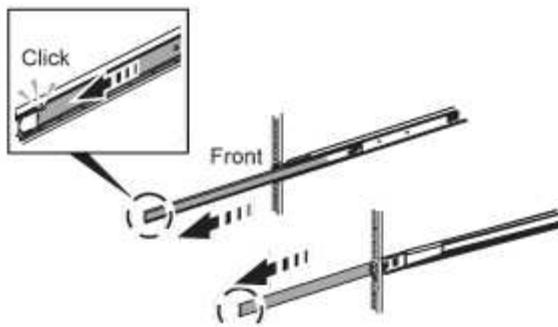


H610S and H615C node/chassis

In the case of H615C and H610S, the terms "node" and "chassis" are used interchangeably because node and chassis are not separate components, unlike in the case of the 2U, four-node chassis.

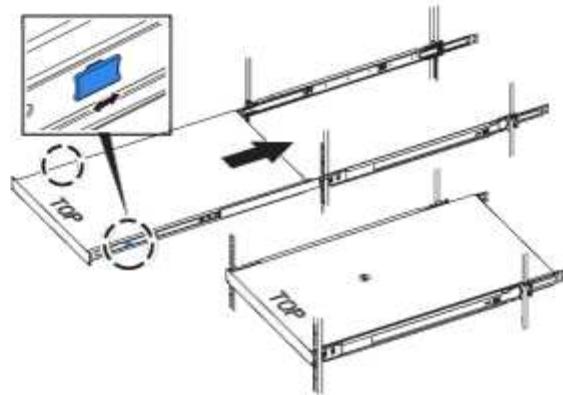
Here is an illustration for installing the node/chassis in the rack:

- i. Extend the middle rail fully toward you.



- ii. Insert the node.

Note: When the node stops moving further, pull the blue latches, one on each side of the node to slide the node all the way in.



Install the switches

If you want to use Mellanox SN2010, SN2100, and SN2700 switches in your NetApp HCI installation, follow the instructions provided here to install and cable the switches:

- [Mellanox hardware user manual](#)
- [TR-4836: NetApp HCI with Mellanox SN2100 and SN2700 Switch Cabling Guide \(login required\)](#)

Cable the nodes

If you are adding nodes to an existing NetApp HCI installation, ensure that the cabling and network configuration of the nodes that you add are identical to the existing installation.



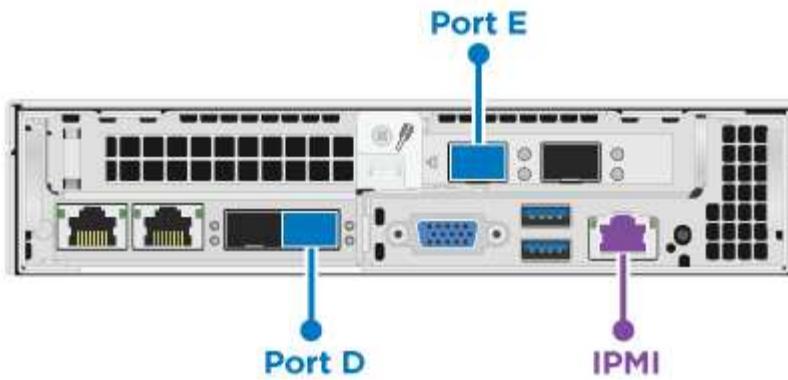
Ensure that the airflow vents at the rear of the chassis are not blocked by cables or labels. This can lead to premature component failures due to overheating.

- [H410C compute node and H410S storage node](#)
- [H610C compute node](#)
- [H615C compute node](#)
- [H610S storage node](#)

H410C compute node and H410S storage node

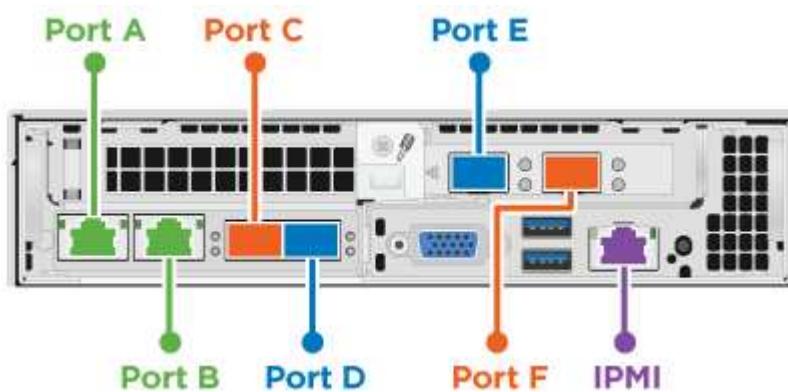
You have two options for cabling the H410C node: using two cables or using six cables.

Here is the two-cable configuration:



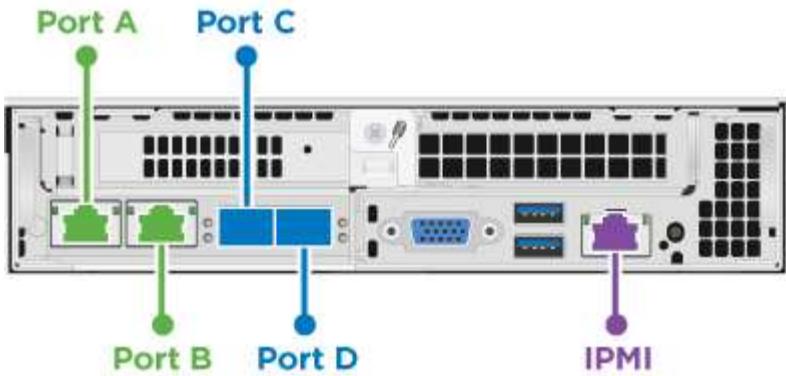
- For ports D and E, connect two SFP28/SFP+ cables or transceivers for shared management, virtual machines, and storage connectivity.
- (Optional, recommended) Connect a CAT5e cable in the IPMI port for out-of-band management connectivity.

Here is the six-cable configuration:



- For ports A and B, connect two CAT5e or higher cables in ports A and B for management connectivity.
- For ports C and F, connect two SFP28/SFP+ cables or transceivers for virtual machine connectivity.
- For ports D and E, connect two SFP28/SFP+ cables or transceivers for storage connectivity.
- (Optional, recommended) Connect a CAT5e cable in the IPMI port for out-of-band management connectivity.

Here is the cabling for the H410S node:



- For ports A and B, connect two CAT5e or higher cables in ports A and B for management connectivity.
- For ports C and D, connect two SFP28/SFP+ cables or transceivers for storage connectivity.
- (Optional, recommended) Connect a CAT5e cable in the IPMI port for out-of-band management connectivity.

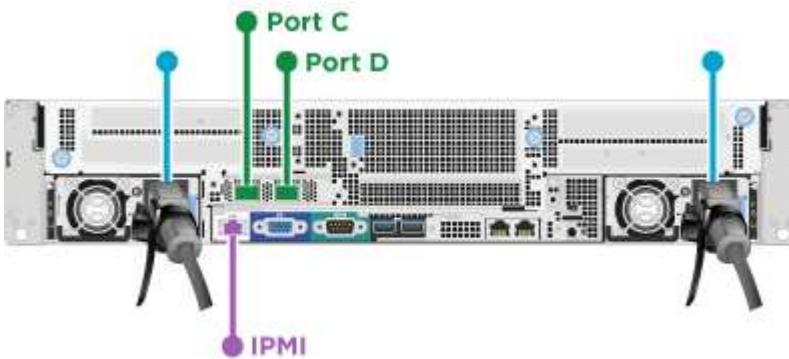
After you cable the nodes, connect the power cords to the two power supply units per chassis and plug them into 240V PDU or power outlet.

H610C compute node

Here is the cabling for the H610C node:



H610C nodes are deployed only in the two-cable configuration. Ensure that all the VLANs are present on ports C and D.



- For ports C and D, connect the node to a 10/25GbE network using two SFP28/SFP+ cables.
- (Optional, recommended) Connect the node to a 1GbE network using an RJ45 connector in the IPMI port.
- Connect both power cables to the node, and plug the power cables to a 200-240V power outlet.

H615C compute node

Here is the cabling for the H615C node:



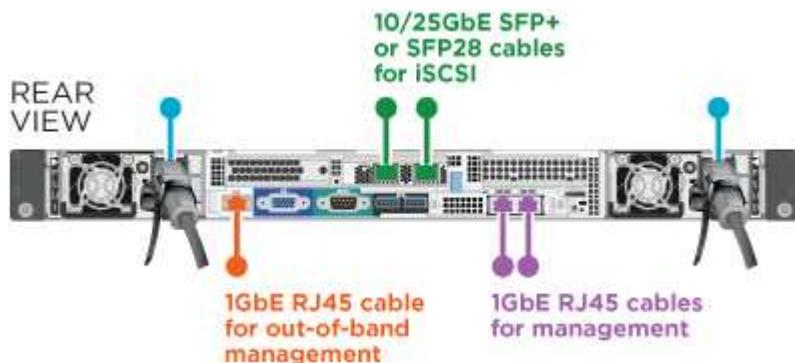
H615C nodes are deployed only in the two-cable configuration. Ensure that all the VLANs are present on ports A and B.



- For ports A and B, connect the node to a 10/25GbE network using two SFP28/SFP+ cables.
- (Optional, recommended) Connect the node to a 1GbE network using an RJ45 connector in the IPMI port.
- Connect both power cables to the node, and plug the power cables to a 110-140V power outlet.

H610S storage node

Here is the cabling for the H610S node:



- Connect the node to a 1GbE network using two RJ45 connector in the IPMI port.
- Connect the node to a 10/25GbE network using two SFP28 or SFP+ cables.
- Connect the node to a 1GbE network using an RJ45 connector in the IPMI port.
- Connect both power cables to the node.

Power on the nodes

It takes approximately six minutes for the nodes to boot.

Here is an illustration that shows the power button on the NetApp HCI 2U chassis:

Here is an illustration that shows the power button on the H610C node:



Here is an illustration that shows the power button on the H615C and H610S nodes:



Configure NetApp HCI

Choose from one of the following options:

- [New NetApp HCI installation](#)
- [Expand an existing NetApp HCI installation](#)

New NetApp HCI installation

Steps

1. Configure an IPv4 address on the management network (Bond1G) on one NetApp HCI storage node.



If you are using DHCP on the management network, you can connect to the DHCP-acquired IPv4 address of the storage system.

- a. Plug in a keyboard, video, mouse (KVM) to the back of one storage node.
- b. Configure the IP address, subnet mask, and gateway address for Bond1G in the user interface. You can also configure a VLAN ID for the Bond1G network.

2. Using a supported web browser (Mozilla Firefox, Google Chrome, or Microsoft Edge), navigate to the NetApp Deployment Engine by connecting to the IPv4 address that you configured in Step 1.

3. Use the NetApp Deployment Engine user interface (UI) to configure NetApp HCI.



All the other NetApp HCI nodes will be discovered automatically.

Expand an existing NetApp HCI installation

Steps

1. Open the IP address of the management node in a web browser.
2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.
3. Follow the steps in the wizard to add storage and/or compute nodes to your NetApp HCI installation.



To add H410C compute nodes, the existing installation must run NetApp HCI 1.4 or later. To add H615C compute nodes, the existing installation must run NetApp HCI 1.7 or later.



The newly installed NetApp HCI nodes on the same network will be discovered automatically.

Perform post-configuration tasks

Depending on the type of node you have, you might need to perform additional steps after you install the hardware and configure NetApp HCI.

- [H610C node](#)
- [H615C and H610S nodes](#)

H610C node

Install the GPU drivers in ESXi for each H610C node that you installed, and validate their functionality.

H615C and H610S nodes

Steps

1. Use a web browser and navigate to the default BMC IP address: 192.168.0.120
2. Log in using user name `root` and password `calvin`.
3. From the node management screen, navigate to **Settings > Network Settings**, and configure the network parameters for the out-of-band management port.

If your H615C node has GPUs in it, install GPU drivers in ESXi for each H615C node that you installed, and validate their functionality.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [TR-4820: NetApp HCI Networking Quick Planning Guide](#)
- [NetApp Configuration Advisor 5.8.1 or later network validation tool](#)

Configure LACP for optimal storage performance

For optimal NetApp HCI storage cluster performance, you should configure Link Aggregation Control Protocol (LACP) on the switch ports used for each of the storage nodes.

What you'll need

- You have configured the switch ports connected to the 10/25GbE interfaces of NetApp HCI storage nodes as LACP port channels.
- You have set the LACP timers on the switches handling storage traffic to “fast mode (1s)” for optimal failover detection time. During deployment, the Bond1G interfaces on all storage nodes are automatically configured for active/passive mode.
- You have configured Cisco Virtual PortChannel (vPC) or the equivalent switch stacking technology for the switches handling the storage network. Switch stacking technology eases configuration of LACP and port channels, and provides a loop-free topology between switches and the 10/25GbE ports on the storage nodes.

Steps

1. Follow your switch vendor recommendations for enabling LACP on the switch ports used for NetApp H-series storage nodes.
2. Change the bond mode on all storage nodes to LACP in the on-node user interface (also known as the terminal user interface, or TUI) before you deploy NetApp HCI.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Validate your environment with Active IQ Config Advisor

Before you rack NetApp HCI hardware and perform the installation of NetApp HCI, you need to verify that your environment meets NetApp HCI networking requirements. Active IQ Config Advisor runs checks on your environment by validating network, switch, and VMware vSphere configurations. The tool generates a report which you can use to help you resolve issues, and you can forward the report to your Professional Services engineer to prepare and schedule an installation.

Install Active IQ Config Advisor

Download and install Active IQ Config Advisor on a PC that has access to the NetApp HCI networks.

Steps

1. In a web browser, select **Tools** from the NetApp Support menu, search for Active IQ Config Advisor, and download the tool.

[NetApp Support Site > Tools](#).

After you agree to the End User License Agreement (EULA), the Download page appears. Microsoft Windows, Linux, and Mac binaries are available in the **Client Tool** pane.

2. Run the executable.
3. Select a language, and select **OK**.
4. Select **Next**.
5. Read the EULA and select **I Agree**.
6. Select **Install**.
7. Ensure that **Run Active IQ Config Advisor** is selected, and select **Finish**.

After a short delay, the Active IQ Config Advisor UI opens in a new browser window or tab.

Use Active IQ Config Advisor

Active IQ Config Advisor runs in a browser window, collects information about your network and environment, and generates a report you can use to resolve any network or configuration issues that might interfere with NetApp HCI deployment.

What you'll need

You have installed Active IQ Config Advisor on a device that can access the management network, VMware vCenter Server networking (if you are joining an existing VMware installation), and switches that will be used for NetApp HCI.



If you are using Mellanox switches and NetApp Professional Services is configuring them as part of deployment, you do not need to provide switch information.

About this task

Active IQ Config Advisor performs only read-only checks to gather information. No configuration is modified as part of the collection.

Steps

1. Open Active IQ Config Advisor.

Config Advisor appears with the **Basic Settings** window in a web browser. Here, you can define global collection settings and encrypt the collection results.

2. Enter a passphrase in the **Encryption Settings** section to encrypt the collection project.

This ensures that only you are able to load this collection project after it is created.

3. Identify this collection report as yours by entering your name and email address in the **User Verification** section.
4. Select **Save**.
5. Select **Create a new data collection**.
6. Select **Solution Based** in the **Collection Type** drop-down menu.
7. Select **NetApp HCI Pre Deployment** in the **Profile** drop-down menu.
8. For each type of device in the **Type** column, select the number of that type of device in your NetApp HCI network in the **Actions** drop-down menu.

For example, if you have three Cisco switches, choose 3 from the **Actions** column drop-down menu in that row. Three rows appear, one for each Cisco switch you identified.



If you are using Mellanox switches and NetApp Professional Services is configuring them as part of deployment, you do not need to provide switch information.

9. For any switches that you identified, enter the management IP address and administrator credentials.
10. For any VMware vCenter Servers you identified, do one of the following:
 - If you are deploying a new vCenter Server, provide the IP address or Fully Qualified Domain Name (FQDN) that is planned for the server.
 - If you are joining an existing vCenter Server, provide the IP address or FQDN and the administrator credentials for the server.
11. Optional: If you added information for switches, enter the number of compute and storage nodes in the **Switch Validation** section.
12. Choose which compute node cabling configuration you plan to use in the **Compute node network** section.
13. Enter individual switch ports and any VLAN tags you plan to use for the management, vMotion, and storage networks for any switches in the **Compute node network** section.
14. Enter individual switch ports and any VLAN tags you plan to use for the management and storage networks for any switches in the **Storage node network** section.
15. In the **Network Settings Check** section, enter the IP addresses and gateway IP address for the management network, followed by lists of servers for DNS, NTP, and vCenter Server (if you are deploying a new vCenter Server with NetApp HCI).

This section enables Active IQ Config Advisor to ensure that the management network is available for use, and also ensures that services such as DNS and NTP are working properly.

16. Select **Validate** to ensure all of the IP address information and credentials you have entered are valid.
17. Select **Save or Collect**.

This starts the collection process, and you can see the progress as the collection runs along with a real-time log of the collection commands. The **Progress** column shows color-coded progress bars for each collection task.



The progress bars use the following colors to show status:

- **Green:** The collection has finished with no command failures. You can see the deployment risks and recommendations by selecting the **View & Analyze** icon in the **Actions** menu.
 - **Yellow:** The collection has finished with some command failures. You can see the deployment risks and recommendations by selecting the **View & Analyze** icon in the **Actions** menu.
 - **Red:** The collection has failed. You need to resolve the errors and run the collection again.
18. Optional: When the collection is complete, you can select the binocular icon for any collection row to see the commands that were run and the data that was collected.
 19. Select the **View & Analyze** tab.

This page shows a general health report of your environment. You can select a section of the pie chart to see more details about those specific checks or descriptions of problems, along with recommendations on resolving any issues that might interfere with successful deployment. You can resolve these issues yourself or request help from NetApp Professional Services.

20. Select **Export** to export the collection report as a PDF or Microsoft Word document.



PDF and Microsoft Word document outputs include the switch configuration information for your deployment, which NetApp Professional Services uses to verify the network settings.

21. Send the exported report file to your NetApp Professional Services representative.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Configure IPMI for each node

After you have racked, cabled, and powered on your NetApp HCI hardware, you can configure Intelligent Platform Management Interface (IPMI) access for each node. Assign each IPMI port an IP address and change the default IPMI administrator password as soon as you have remote IPMI access to the node.

Prerequisites

After you have validated that your environment is ready to support NetApp HCI and resolved any potential issues, you need to complete some final tasks before deployment.

- Ensure you have a successful report from Active IQ Config Advisor.
- Gather all relevant information about your network, current or planned VMware infrastructure, and planned user credentials.
- Rack, cable, and power on the NetApp HCI installation.

Manually assign the IPMI port IP address

Dynamic Host Configuration Protocol (DHCP) is enabled by default for the IPMI port of each NetApp HCI node. If your IPMI network does not use DHCP, you can manually assign a static IPv4 address to the IPMI port.

What you'll need

Ensure that you have a keyboard, video, and mouse (KVM) switch or monitor and keyboard you can use to access the BIOS of each node.

About this task

Use the arrow keys to navigate in the BIOS. Select a tab or option by pressing `Enter`. Go back to previous screens by pressing `ESC`.

Steps

1. Power on the node.
2. Upon booting, enter the BIOS by pressing the `Del` key.
3. Select the IPMI tab.
4. Select **BMC Network Configuration** and press `Enter`.
5. Choose **Yes** and press `Enter`.

6. Select **Configuration Address Source** and press **Enter**.
7. Choose **Static** and press **Enter**.
8. Select **Station IP address** and enter a new IP address for the IPMI port. Press **Enter** when finished.
9. Select **Subnet mask** and enter a new subnet mask for the IPMI port. Press **Enter** when finished.
10. Select **Gateway IP address** and enter a new gateway IP address for the IPMI port. Press **Enter** when finished.
11. Connect one end of an Ethernet cable to the IPMI port and the other end to a switch.

The IPMI port for this node is ready to use.

12. Repeat this procedure for any other NetApp HCI nodes with IPMI ports that are not configured.

Change the default IPMI password for H410C and H410S nodes

You should change the default password for the IPMI administrator account on each compute and storage node as soon as you configure the IPMI network port.

What you'll need

You have configured the IPMI IP address for each compute and storage node.

Steps

1. Open a web browser on a computer that can reach the IPMI network and browse to the IPMI IP address for the node.
2. Enter the user name **ADMIN** and password **ADMIN** in the login prompt.
3. Upon logging in, select the **Configuration** tab.
4. Select **Users**.
5. Select the **ADMIN** user and select **Modify User**.
6. Select the **Change Password** check box.
7. Enter a new password in the **Password** and **Confirm Password** fields.
8. Select **Modify**, and then select **OK**.
9. Repeat this procedure for any other NetApp HCI H410C and H410S nodes with default IPMI passwords.

Change the default IPMI password for H610C, H615C, and H610S nodes

You should change the default password for the IPMI administrator account on each compute and storage node as soon as you configure the IPMI network port.

What you'll need

You have configured the IPMI IP address for each compute and storage node.

Steps

1. Open a web browser on a computer that can reach the IPMI network and browse to the IPMI IP address for the node.
2. Enter the user name **root** and password **calvin** in the login prompt.
3. Upon logging in, select the menu navigation icon at the top left of the page to open the sidebar drawer.

4. Select **Settings**.
5. Select **User Management**.
6. Select the **Administrator** user from the list.
7. Enable the **Change Password** check box.
8. Enter a new, strong password in the **Password** and **Confirm Password** fields.
9. Select **Save** at the bottom of the page.
10. Repeat this procedure for any other NetApp HCI H610C, H615C, or H610S nodes with default IPMI passwords.

Find more information

- [NetApp SolidFire Active IQ Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

Deploy NetApp HCI

Access the NetApp Deployment Engine

Access the NetApp Deployment Engine

To deploy NetApp HCI, you need to access the NetApp Deployment Engine on one of the NetApp H-Series storage nodes via the IPv4 address assigned to the Bond1G interface, which is the logical interface that combines ports A and B for storage nodes. This storage node becomes the controlling storage node for the deployment process. Depending on your environment, you need to either configure the IPv4 address or retrieve it from one of the storage nodes.



You can only access the NetApp Deployment Engine using the Bond1G interface of a storage node. Using the Bond10G interface, the logical interface that combines ports C and D for storage nodes, is not supported.

Use one of the following methods that best describes your network environment to access the NetApp Deployment Engine:

Scenario	Method
You do not have DHCP in your environment	Access the NetApp Deployment Engine in environments without DHCP
You have DHCP in your environment	Access the NetApp Deployment Engine in environments with DHCP
You want to assign all IP addresses manually	Manually assign IP addresses to access the NetApp Deployment Engine

Find more information

- [Configure Fully Qualified Domain Name web UI access](#)

Access the NetApp Deployment Engine in environments without DHCP

When DHCP is not in use on the network, you need to set a static IPv4 address on the Bond1G interface of one of the storage nodes (also known as a controlling storage node) that you will use to access the NetApp Deployment Engine. The NetApp Deployment Engine on the controlling storage node will discover and communicate with other compute and storage nodes using IPv4 addresses that have been auto-configured on the Bond10G interfaces of all nodes. You should use this method unless your network has special requirements.

What you'll need

- You or your network administrator have completed the tasks in the Installation and Setup Instructions document.
- You have physical access to the NetApp HCI nodes.

- All of the NetApp HCI nodes are powered on.
- DHCP is not enabled for the NetApp HCI networks and the NetApp HCI nodes have not obtained IP addresses from DHCP servers.
- The NetApp HCI management network is configured as the native VLAN on the Bond1G and Bond10G interfaces of all nodes.

Steps

1. Plug a KVM into the back of one of the NetApp HCI storage nodes (this node will become the controlling storage node).
2. Configure the IP address, subnet mask, and gateway address for Bond1G in the user interface. You can also configure a VLAN ID for the Bond1G network if needed.



You cannot reuse this IPv4 address later during deployment with the NetApp Deployment Engine.

3. Open a web browser on a computer that can access the NetApp HCI management network.
4. Browse to the IP address you assigned to the controlling storage node. For example:

```
http://<Bond1G IP address>
```

This takes you to the NetApp Deployment Engine user interface.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI](#) and [firmware versions for NetApp HCI storage nodes](#)

Access the NetApp Deployment Engine in environments with DHCP

In environments where servers automatically acquire IPv4 configuration from DHCP, you can access the NetApp Deployment Engine using the IPv4 address assigned to the Bond1G interface on one of the storage nodes. You can use a USB stick to retrieve the IPv4 address from one of the storage nodes. The NetApp Deployment Engine will automatically discover other compute and storage nodes that use DHCP-assigned IPv4 addresses. You should not use this method unless your network has special requirements.

What you'll need

- You or your network administrator have completed the tasks in the Installation and Setup Instructions document.
- You have physical access to the NetApp HCI nodes.
- All of the NetApp HCI nodes are powered on.
- DHCP is enabled on the NetApp HCI management and storage networks.
- The DHCP address pool is large enough to accommodate two IPv4 addresses per NetApp HCI node.



For the NetApp HCI deployment to succeed, all nodes in the deployment must either have DHCP-acquired or auto-configured IPv4 addresses (you cannot mix IPv4 address assignment methods).

About this task

If DHCP is in use only for the storage network (Bond10G interfaces), you should use the steps outlined in [xref:./docs/Access the NetApp Deployment Engine in environments without DHCP](#) to access the NetApp Deployment Engine.

Steps

1. Wait several minutes for the nodes to request IP addresses.
2. Choose a storage node and insert a USB stick into the node. Leave it in for at least five seconds.
3. Remove the USB stick, and insert it into your computer.
4. Open the `readme.html` file. This takes you to the NetApp Deployment Engine user interface.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI](#) and [firmware versions for NetApp HCI storage nodes](#)

Manually assign IP addresses to access the NetApp Deployment Engine

You can manually assign static IPv4 addresses to the Bond1G and Bond10G interfaces on all NetApp HCI nodes to access the NetApp Deployment Engine and deploy NetApp HCI. You should not use this method unless your network has special requirements.

What you'll need

- You or your network administrator have completed the tasks in the Installation and Setup Instructions document.

- You have physical access to the NetApp HCI nodes.
- All of the NetApp HCI nodes are powered on.

- DHCP is not enabled for the NetApp HCI networks and the NetApp HCI nodes have not obtained IP addresses from DHCP servers.

NOTE: All IP addresses you assign manually before using the NetApp Deployment Engine to deploy the system are temporary and cannot be reused. If you choose to manually assign IP addresses, you need to set aside a second permanent set of unused IP addresses that you can assign during final deployment.

About this task

In this configuration, compute and storage nodes will use static IPv4 addresses to discover and communicate with other nodes during deployment. This configuration is not recommended.

Steps

1. Plug a KVM into the back of one of the NetApp HCI storage nodes (this node will become the controlling storage node).
2. Configure the IP address, subnet mask, and gateway address for Bond1G and Bond10G in the user interface. You can also configure a VLAN ID for each network if needed.
3. Repeat step 2 for the remaining storage and compute nodes.

4. Open a web browser on a computer that can access the NetApp HCI management network.
5. Browse to the Bond1G IP address you assigned to the controlling storage node. For example:

```
http://<Bond1G IP address>
```

This takes you to the NetApp Deployment Engine user interface.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)

Start your deployment

Before continuing with your NetApp HCI deployment, you need to read and understand the end user license agreements.

Steps

1. On the **Welcome to NetApp HCI** page, select **Get Started**.
2. On the **Prerequisites** page, do the following:
 - a. Ensure each prerequisite is met, and select each associated checkbox to confirm.
 - b. Select **Continue**.
3. On the **End User Licenses** page, do the following:
 - a. Read the NetApp End User License Agreement
 - b. If you accept the terms, select **I accept** at the bottom of the agreement text.
 - c. Read the VMware End User License Agreement.
 - d. If you accept the terms, select **I accept** at the bottom of the agreement text.
 - e. Select **Continue**.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)

Configure VMware vSphere

VMware vSphere configuration

NetApp HCI uses the vCenter Server and ESXi components of VMware vSphere. vCenter Server is used to manage and monitor the VMware ESXi hypervisor installed on each compute node. You can install and configure a new vSphere deployment, which also installs the NetApp Element Plug-in for vCenter Server, or you can join and extend an existing vSphere deployment.

Be aware of the following caveats when you use the NetApp Deployment Engine to install a new vSphere deployment:

- The NetApp Deployment Engine installs the new vCenter Server Appliance with the Small deployment size option.
- The vCenter Server license is a temporary evaluation license. For continued operation after the evaluation period, you need to obtain a new license key from VMware and add it to the vCenter Server license inventory.



If your vSphere inventory configuration uses a folder to store the NetApp HCI cluster within the vCenter datacenter, some operations, such as expanding NetApp HCI compute resources, will fail. Ensure that the NetApp HCI cluster is directly under the datacenter in the vSphere web client inventory tree, and is not stored in a folder. See the NetApp Knowledgebase article for more information.

If you install a new vCenter Server, you can install a vSphere standard switch or a vSphere distributed switch (VDS) during network configuration. A VDS enables a simplified, centralized management of virtual machine network configuration after NetApp HCI deployment. Cloud data services functionality on NetApp HCI requires a VDS; vSphere standard switches are not supported for cloud data services.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)

Configure a new VMware vSphere environment

You can deploy a new vSphere environment as part of the NetApp HCI installation process by providing some of the network information that vSphere should use. Note that if you configure vSphere using an IP address, the address cannot be changed after installation.

What you'll need

You have obtained the network information for the planned vSphere environment.

Steps

1. Select **Configure a new vSphere deployment**.
2. Select which version of vSphere the system should install during deployment.
3. Configure the new vSphere environment using one of the following options:

Option	Steps
Use a domain name (recommended).	<ol style="list-style-type: none"> <li data-bbox="861 160 1421 223">a. Select Configure Using a Fully Qualified Domain Name. <li data-bbox="861 244 1486 350">b. Enter the vCenter Server domain name in the vCenter Server Fully Qualified Domain Name field. <li data-bbox="861 371 1449 435">c. Enter the DNS server IP address in the DNS Server IP Address field. <li data-bbox="861 456 1111 483">d. Select Continue.
Use an IP address.	<ol style="list-style-type: none"> <li data-bbox="861 530 1393 557">a. Select Configure Using an IP Address. <li data-bbox="861 578 1111 606">b. Select Continue.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI](#) and [firmware versions for NetApp HCI storage nodes](#)

Join an existing VMware vSphere deployment

Configure NetApp HCI to take advantage of an existing vSphere deployment by providing the vCenter Server network information and credentials.

What you'll need

- If you are joining an existing vSphere 6.7 deployment, make sure vCenter Server is running version 6.7 Update 1.
- If you are joining an existing vSphere 6.5 deployment, make sure vCenter Server is running version 6.5 Update 2 or later.
- Obtain the network details and administrator credentials for your existing vSphere deployment.

About this task

If you join multiple vCenter Server systems that are connected using vCenter Linked Mode, NetApp HCI only recognizes one of the vCenter Server systems.



Using the NetApp Element Plug-in for vCenter Server to manage cluster resources from other vCenter Servers using [vCenter Linked Mode](#) is limited to local storage clusters only.

Steps

1. Select **Join and extend an existing vSphere deployment**.
2. Enter the domain name or IP address in the **vCenter Server Domain Name or IP address** field. If you enter a domain name, you also need to enter the IP address of an active DNS server in the **DNS Server IP Address** field that appears.
3. Enter the credentials of a vSphere administrator in the **User Name and Password** fields.
4. Select **Continue**.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)

Configuring NetApp HCI credentials

During deployment, you define a set of credentials to be used across the newly deployed VMware vSphere environment, the NetApp HCI compute and storage resources, and the management node. If you are deploying NetApp HCI into an existing vSphere environment, these credentials are not applied to the existing vCenter Server.

About this task

Remember the following points about the credentials you set in the NetApp HCI Deployment Engine:

- **NetApp Hybrid Cloud Control (HCC) or Element UI:** To log in to NetApp HCC or the Element user interface upon successful deployment, use the user name and password specified in this deployment step.
- **VMware vCenter:** To log in to vCenter (if installed as part of deployment), use user name with the suffix `@vsphere.local` or the built-in `Administrator@vsphere.local` user account, and the password specified in this deployment step.
- **VMware ESXi:** To log in to ESXi on the compute nodes, use the user name `root` and the same password specified in this deployment step.

For interaction with VMware vCenter instances, NetApp Hybrid Cloud Control will use one of the following:

- The built-in `Administrator@vsphere.local` user account on the vCenter instance that was installed as part of the deployment.
- The vCenter credentials that were used to connect the NetApp HCI deployment to an existing VMware vCenter Server.

Steps

1. On the **Credentials** page, enter a user name in the **User Name** field.
2. Enter a password in the **Password** field. The password must conform to the password criteria visible in the **Password must contain** box.
3. Confirm the password in the **Re-enter Password** field.
4. Select **Continue**.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)
- To update vCenter and ESXi credentials later, see [Update vCenter or ESXi credentials](#).

Select a network topology

When cabling NetApp HCI nodes, you have the option of using different network cable configurations depending on your needs. For each compute node, you can use all six network ports, with different types of traffic assigned to each pair of ports, or you can use

two ports with all types of traffic assigned to the ports. Storage nodes use the standard four-cable configuration. Your choice affects which compute nodes are selectable in the inventory.

What you'll need

If you choose the two-cable network topology for compute nodes, consider the following requirements:

- You have a VMware vSphere Enterprise Plus license ready to apply after deployment is complete.
- You have verified that the configuration of your network and network switches is correct.
- VLAN tagging is required for storage and vMotion networks for all compute and storage nodes.

Steps

1. On the **Network Topology** page, select a compute node topology that fits the way you installed compute nodes for NetApp HCI:
 - **6 Cable Option:** The six-cable option provides dedicated ports for each type of traffic (management, virtual machine, and storage). You can optionally enable vSphere Distributed Switch (VDS). Enabling VDS configures a distributed switch, enabling simplified, centralized management of virtual machine network configuration after NetApp HCI deployment is complete. If you enable it, you must have a vSphere Enterprise Plus license ready to apply after deployment.
 - **2 Cable Option:** The two-cable option combines management, virtual machine, and storage traffic on two bonded ports. This cabling option requires VDS, and automatically enables it. You must have a vSphere Enterprise Plus license ready to apply after deployment.
2. Some cabling options display multiple back panel views of different types of node hardware. Cycle through the back panel views to see how to connect the network cables for that specific node model and cabling option.
3. When finished, select **Continue**.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI](#) and [firmware versions for NetApp HCI storage nodes](#)

Inventory selection

Inventory selection and node compatibility

When choosing nodes for your deployment, some restrictions apply to the node configurations you can combine in the same deployment.

Storage node compatibility

NetApp HCI supports storage nodes and drives with SED (Self-encrypting drive) and FIPS 140-2 drive encryption capability. When deploying or expanding NetApp HCI, you can mix nodes with different reported levels of encryption, but NetApp HCI only supports the more basic form of encryption in this situation. For example, if you mix a storage node that is FIPS encryption capable with nodes that only support SED encryption, SED encryption is supported with this configuration, but FIPS drive encryption is not.



Adding storage nodes capable of FIPS drive encryption to the storage cluster does not automatically enable the FIPS drive encryption feature. After you deploy or expand an installation with FIPS-capable nodes, you need to manually enable FIPS drive encryption. See the Element Software User Guide for instructions.

All storage nodes must run the same minor version of Element software to be compatible in the same deployment. For example, you cannot mix a storage node running Element 11.3.1 with other storage nodes running Element 11.5.



Depending on node hardware configuration, H410S storage nodes might appear in the inventory list labeled as H300S, H500S, or H700S storage nodes.

NetApp HCI supports only certain storage node models in two-node storage clusters. For more information, see [two-node storage clusters](#) or the Release Notes for your NetApp HCI version.



For two-node storage cluster deployments, the storage node types are limited to nodes with 480GB and 960GB drives.

Compute node compatibility

Compute nodes must meet the following requirements to be selectable as inventory:

- The CPU generations in all compute nodes must match for proper VMware vMotion functionality. After you select a compute node from the inventory, you cannot select compute nodes with different CPU generations.
- You cannot intermix compute nodes with GPU-enabled compute nodes in the same compute cluster. If you select a GPU-enabled compute node, CPU-only compute nodes become unselectable, and vice versa.
- The software version running on the compute node must match the major and minor version of the NetApp Deployment Engine hosting the deployment. If this is not the case, you need to reimage the compute node using the RTFI process. See the NetApp Knowledgebase articles regarding RTFI for instructions.
- The compute node must have the cabling configuration you selected on the Network Topology page to be selectable in the **Compute Nodes** list.
- The network cabling configurations for compute nodes of the same model must match within a single compute cluster.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [SolidFire and Element Software Documentation Center](#)

Select inventory

On the **Inventory** page, the NetApp Deployment Engine automatically detects available compute and storage nodes, enabling you to select and add all NetApp HCI resources to the deployment. If a node does not meet the requirements for deployment, it is not selectable and problems are indicated as errors. You can position your cursor over the error in the node's row to see an explanation. When choosing node inventory on the Inventory page, the storage node that is hosting the NetApp Deployment Engine is automatically selected, and you cannot deselect it.

What you'll need

Jumbo frames must be enabled for proper inventory detection. If no nodes or only a subset of nodes appear in the inventory, verify that the switch ports used for NetApp HCI nodes (all SFP+/SFP28 interfaces) are configured with jumbo frames.

Steps

1. On the **Inventory** page, view the list of available nodes.

If the system cannot detect any inventory, it displays an error. Correct the error before continuing.

If your system uses DHCP for IP address assignment, the storage and compute resources might not appear in the inventory immediately.

2. Optional: If a resource does not appear in the inventory immediately, or if you address an error and need to refresh the inventory, select **Refresh Inventory**. You might need to refresh the inventory multiple times.
3. Optional: To filter the inventory on node attributes, such as node type:
 - a. Select **Filter** in the header of the **Compute Nodes** or **Storage Nodes** lists.
 - b. Choose criteria from the drop-down lists.
 - c. Below the drop-down lists, enter information to satisfy the criteria.
 - d. Select **Add Filter**.
 - e. Clear individual filters by selecting **X** next to an active filter, or clear all filters by selecting **X** above the list of filters.
4. Select all compute nodes that shipped with your system from the **Compute Nodes** list.

You need to select at least two compute nodes to proceed with deployment.

5. Select all storage nodes that shipped with your system from the **Storage Nodes** list.

You need to select at least two storage nodes to proceed with deployment.

6. Optional: If a storage node selection box is flagged, that storage node exceeds 33% of the total storage cluster capacity. Do one of the following:
 - Clear the selection box for the flagged storage node.
 - Select additional storage nodes to more equally distribute the storage cluster capacity between nodes.
7. Select **Continue**.

Find more information

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

Configure network settings

NetApp HCI provides a network settings page with an easy form to simplify network configuration. When you complete the easy form, NetApp HCI automatically populates much of the rest of the information on the network settings page. You can then enter final network settings and verify that the network configuration is correct before proceeding.

You do not need to complete the form in its entirety.

What you'll need

- You have obtained the following information:
 - The planned naming prefix for the hosts and storage cluster
 - All planned subnet mask, starting IP address, default gateway, and VLAN IDs for the management, iSCSI, and vMotion networks
 - The IP address, default gateway, VLAN IDs, and subnet mask information for any planned VMware vCenter deployment
 - The Network Time Protocol (NTP) server address for NetApp HCI
 - The DNS server IP address information for NetApp HCI
- If you are deploying a vSphere Distributed Switch, you have a vSphere Enterprise Plus license ready to apply after deployment is complete.
- If you assigned VLAN IDs to node ports during terminal user interface (TUI) configuration, you have configured those ports with the same VLAN ID during network configuration. You do not need to configure tagged host ports as access ports or native VLANs on the connected switch ports.
- You have verified that your network switch configuration is correct. Incorrect switch configurations (such as incorrect VLANs or MTU size) will cause deployment errors.

About this task

If you selected the two-cable network topology for compute nodes, you must use VLAN IDs for the vMotion and storage networks for all compute and storage nodes in the deployment (VLAN IDs are optional for the management networks).



In environments that require host-side VLAN tagging before deployment, if you have configured VLAN IDs on compute and storage nodes so they are discoverable by the NetApp Deployment Engine, ensure you use the correct VLANs when configuring network settings in the NetApp Deployment Engine.

If you are deploying using a two-node or three-node storage cluster, you can complete IP address information for Witness Nodes on the **Network Settings** page after using the easy form.

Steps

1. Optional: To disable live validation of network information you enter on this page, toggle the **Live network validation is** switch to **Off**.
2. In the **Infrastructure Services** section of the **Network Settings** page, enter the DNS and NTP server information for NetApp HCI in the following fields:

Field	Description
DNS Server IP Address 1	The IP address of the primary DNS server for NetApp HCI. If you specified a DNS server on the vCenter Configuration page, this field is populated and read-only.
DNS Server IP Address 2 (Optional)	An optional IP address of a secondary DNS server for NetApp HCI.
NTP Server Address 1	The IP address or fully qualified domain name of the primary NTP server for this infrastructure.

Field	Description
NTP Server Address 2 (Optional)	An optional IP address or fully qualified domain name of the secondary NTP server for this infrastructure.

3. Select **To save time, launch the easy form** to enter fewer network settings.

The **Network Settings Easy Form** dialog box appears.

4. Enter a naming prefix in the **Naming Prefix** field.

The naming prefix is applied to the name of each host and the name of the storage cluster. Naming prefixes have the following characteristics:

- Must start with a letter
- Can contain letters, numbers, and hyphens
- Cannot exceed 55 characters

5. Choose one of the following options for assigning VLAN IDs.

At any time when using the form, select **Clear** next to a row to clear input from a row of fields.



When you assign VLAN IDs, you are configuring VLAN tags that NetApp HCI will apply to the network traffic. You do not need to enter your native VLAN as a VLAN ID; to use the native VLAN for a network, leave the appropriate field empty.

Option	Steps
Assign VLAN IDs	<ol style="list-style-type: none"> 1. Select Yes for the Will you assign VLAN IDs option. 2. In the VLAN ID column, enter a VLAN tag to use for each type of network traffic you want to assign to a VLAN. Both compute vMotion traffic and iSCSI traffic must use an unshared VLAN ID. 3. Select Next. 4. In the Subnet column, enter subnet definitions in CIDR format for each type of network traffic in each network; for example, 192.168.1.0/24. 5. In the Default Gateway column, enter the IP address of the default gateway for each type of network traffic in each network. 6. In the Starting IP column, enter the first useable IP address for each network subnet in each network.

Option	Steps
Do not assign VLAN IDs	<ol style="list-style-type: none"> 1. Select No for the Will you assign VLAN IDs option. 2. In the Subnet column, enter subnet definitions in CIDR format for each type of network traffic in each network; for example, 192.168.1.0/24. 3. In the Default Gateway column, enter the IP address of the default gateway for each type of network traffic in each network. 4. In the Starting IP column, enter the first useable IP address for each type of network traffic in each network.

6. Select **Apply to Network Settings**.

7. Select **Yes** to confirm.

This populates the **Network Settings** page with the settings you entered in the easy form. NetApp HCI validates the IP addresses you entered. You can disable this validation with the **Disable Live Network Validation** button.

8. Verify that the automatically populated data is correct.

9. Select **Continue**.

Find more information

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

Review and deploy the configuration

You can review the information you provided before beginning deployment. You can also correct any incorrect or incomplete information before you proceed.



During deployment, the management node installation process creates volumes with names beginning with `NetApp-HCI-` in the Element storage cluster, and a SolidFire account beginning with the name `tenant_`. Do not delete these volumes or accounts; doing so will cause a loss in management functionality.

Steps

1. Optional: Select the **Download** icon to download installation information in CSV format.

You can save this file and refer to it later for configuration information.

2. Expand each section and review the information. To expand all sections at once, select **Expand All**.

3. Optional: To make changes to information in any displayed section:

- a. Select **Edit** in the corresponding section.
 - b. Make the necessary changes.
 - c. Select **Continue** until you reach the **Review** page. Your previous settings are saved on each page.
 - d. Repeat steps 2 and 3 to make any other necessary changes.
4. If you do not want to send cluster statistics and support information to NetApp-hosted SolidFire Active IQ servers, clear the final checkbox.

This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve problems before production is affected.

5. If all information is correct, select **Start Deployment**.

A dialog box appears. In the event of network connectivity issues or power loss during the final setup process, or if your browser session is lost, you can copy the URL displayed in the dialog and use it to browse to the final setup progress page.

6. Review the information in the dialog and select **Copy to Clipboard** to copy the URL to your clipboard.
7. Save the URL to a text file on your computer.
8. When you are ready to proceed with deployment, select **OK**.

Deployment begins and a progress page is displayed. Do not close the browser window or navigate away from the progress page until deployment is complete. If your browser session is lost for any reason, you can browse to the URL you copied earlier (and accept any security warnings that appear) to regain access to the final setup progress page.



If the deployment fails, save any error message text and contact NetApp Support.

After deployment is complete, the compute nodes might reboot more than once before becoming ready for service.

After you finish

Begin using NetApp HCI by selecting **Launch vSphere**.

- For NetApp HCI installations using vSphere 6.7, this link launches the HTML5 vSphere web interface. For installations using vSphere 6.5, this link launches the Adobe Flash vSphere web interface.
 - In two storage or three storage node configurations, the NDE configures the Witness Nodes to use the local datastore on the compute nodes. As a result, your vSphere Client displays two **Datastore usage on disk** warnings. To continue, select the **Reset To Green** link in each warning.

Find more information

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

Post-deployment tasks

Post-deployment tasks

Depending on your choices during the deployment process, you need to complete some final tasks before your NetApp HCI system is ready for production use, such as updating firmware and drivers and making any needed final configuration changes.

- [Supported networking changes](#)
- [Disable the smartd service on NetApp HCI compute nodes](#)
- [Disable the "lACP-individual" command on configured switches](#)
- [Create a NetApp HCC role in vCenter](#)
- [Keep VMware vSphere up to date](#)
- [Install GPU drivers for GPU-enabled compute nodes](#)
- [Configure Fully Qualified Domain Name web UI access](#)
- [Access NetApp Hybrid Cloud Control](#)
- [Reduce boot media wear on a NetApp HCI compute node](#)

Find more information

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

Supported networking changes

After you deploy NetApp HCI, you can make limited changes to the default networking configuration. However, certain settings are required for smooth operation and proper network detection. Changing these settings will cause unexpected behavior, and might prevent you from expanding compute and storage resources.

After you deploy your system, you can make the following changes to the default network configuration in VMware vSphere as dictated by your network requirements:

- Change vSwitch names
- Change port group names
- Add and remove additional port groups
- Change the vmnic interface failover order for any additional port groups you have added

H300E, H500E, H700E and H410C compute nodes

NetApp HCI expects the following network configuration for H300E, H500E, H700E and H410C nodes.

The following is a six-interface configuration with VMware vSphere Distributed Switching (VDS). This configuration is only supported when used with VMware vSphere Distributed Switches, and requires VMware vSphere Enterprise Plus licensing.

Network function	vmkernel	vmnic (physical interface)
Management	vmk0	vmnic2 (Port A), vmnic3 (Port B)
iSCSI-A	vmk1	vmnic5 (Port E)
iSCSI-B	vmk2	vmnic1 (Port D)
vMotion	vmk3	vmnic4 (Port C), vmnic0 (Port F)

The following is a six-interface configuration with VMware vSphere Standard Switching (VSS). This configuration uses VMware vSphere Standard Switches (VSS).

Network function	vmkernel	vmnic (physical interface)
Management	vmk0	vmnic2 (Port A), vmnic3 (Port B)
iSCSI-A	vmk2	vmnic1 (Port E)
iSCSI-B	vmk3	vmnic5 (Port D)
vMotion	vmk1	vmnic4 (Port C), vmnic0 (Port F)

The following is a two-interface configuration. This configuration is only supported when used with VMware vSphere Distributed Switches (VDS), and requires VMware vSphere Enterprise Plus licensing.

Network function	vmkernel	vmnic (physical interface)
Management	vmk0	vmnic1 (Port D), vmnic5 (Port E)
iSCSI-A	vmk1	vmnic1 (Port E)
iSCSI-B	vmk2	vmnic5 (Port D)
vMotion	vmk3	vmnic1 (Port C), vmnic5 (Port F)

H610C compute nodes

NetApp HCI expects the following network configuration for H610C nodes.

This configuration is only supported when used with VMware vSphere Distributed Switches (VDS), and requires VMware vSphere Enterprise Plus licensing.



Ports A and B are unused on the H610C.

Network function	vmkernel	vmnic (physical interface)
Management	vmk0	vmnic2 (Port C), vmnic3 (Port D)
iSCSI-A	vmk1	vmnic3 (Port D)
iSCSI-B	vmk2	vmnic2 (Port C)
vMotion	vmk3	vmnic2 (Port C), vmnic3 (Port D)

H615C compute nodes

NetApp HCI expects the following network configuration for H615C nodes.

This configuration is only supported when used with VMware vSphere Distributed Switches (VDS), and requires VMware vSphere Enterprise Plus licensing.

Network function	vmkernel	vmnic (physical interface)
Management	vmk0	vmnic0 (Port A), vmnic1 (Port B)
iSCSI-A	vmk1	vmnic0 (Port B)
iSCSI-B	vmk2	vmnic1 (Port A)
vMotion	vmk3	vmnic0 (Port A), vmnic1 (Port B)

Find more information

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

Disable the smartd service on NetApp HCI compute nodes

By default, the `smartd` service periodically polls the drives in your compute nodes. You should disable this service on all compute nodes after you deploy NetApp HCI.

Steps

1. Using SSH or a local console session, log in to VMware ESXi on the compute node using root credentials.
2. Stop the running `smartd` service:

```
/etc/init.d/smartd stop
```

3. Prevent the `smartd` service from starting at boot:

```
chkconfig smartd off
```

4. Repeat these steps on the rest of the compute nodes in your installation.

Find more information

- [Turn off the smartd service in VMware ESXi](#)
- [VMware KB article 2133286](#)

Disable the "lacp-individual" command on configured switches

By default, the Mellanox switch `lacp-individual` command and the Cisco switch `lacp suspend-individual` command remain configured post deployment. This command is not required post installation; if it remains configured, it can cause volume access issues when troubleshooting or rebooting a switch. Post deployment, you should

check each Mellanox switch and Cisco switch configuration and remove the lacp-individual or lacp suspend-individual command.

Steps

1. Using SSH, open a session to the switch.
2. Show the running configuration:

```
show running-config
```

3. Check the switch configuration output for the lacp-individual or lacp suspend-individual command.



The xxx-xxx is your user supplied interface number(s). If required, you can access the interface number by displaying the Multi-chassis Link Aggregation Group interfaces: show mlag interfaces

- a. For a Mellanox switch, check if the output contains the following line:

```
interface mlag-port-channel xxx-xxx lacp-individual enable force
```

- b. For a Cisco switch, check if the output contains the following line:

```
interface mlag-port-channel xxx-xxx lacp suspend-individual enable force
```

4. If the command is present, remove it from the configuration.

- a. For a Mellanox switch:

```
no interface mlag-port-channel xxx-xxx lacp-individual enable force
```

- b. For a Cisco switch:

```
no interface mlag-port-channel xxx-xxx lacp suspend-individual enable force
```

5. Repeat these steps for each switch in your configuration.

Find more information

- [Storage node goes down during troubleshooting](#)

Keep VMware vSphere up to date

After deploying NetApp HCI, you should use VMware vSphere Lifecycle Manager to apply the latest security patches for the version of VMware vSphere used with NetApp HCI.

Use the [Interoperability Matrix Tool](#) to ensure that all versions of software are compatible. See the [VMware vSphere Lifecycle Manager documentation](#) for more information.

Find more information

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

- [SolidFire and Element Software Documentation Center](#)

Install GPU drivers for GPU-enabled compute nodes

Compute nodes with NVIDIA graphics processing units (GPUs), like the H610C, need NVIDIA software drivers installed in VMware ESXi so that they can take advantage of the increased processing power. After deploying compute nodes with GPUs, you need to perform these steps on each GPU-enabled compute node to install the GPU drivers in ESXi.

Steps

1. Open a browser and browse to the NVIDIA licensing portal at the following URL:

```
https://nvid.nvidia.com/dashboard/
```

2. Download one of the following driver packages to your computer, depending on your environment:

vSphere version	Driver package
vSphere 6.5	NVIDIA-GRID-vSphere-6.5-410.92-410.91-412.16.zip
vSphere 6.7	NVIDIA-GRID-vSphere-6.7-410.92-410.91-412.16.zip

3. Extract the driver package on your computer.

The resulting .VIB file is the uncompressed driver file.

4. Copy the .VIB driver file from your computer to ESXi running on the compute node. The following example commands for each version assume that the driver is located in the \$HOME/NVIDIA/ESX6.x/ directory on the management host. The SCP utility is readily available in most Linux distributions, or available as a downloadable utility for all versions of Windows:

ESXi version	Description
ESXi 6.5	scp \$HOME/NVIDIA/ESX6.5/NVIDIA**.vib root@<ESXi_IP_ADDR>: .
ESXi 6.7	scp \$HOME/NVIDIA/ESX6.7/NVIDIA**.vib root@<ESXi_IP_ADDR>: .

5. Use the following steps to log in as root to the ESXi host and install the NVIDIA vGPU Manager in ESXi.

- a. Run the following command to log in to the ESXi host as the root user:

```
ssh root@<ESXi_IP_ADDRESS>
```

- b. Run the following command to verify that no NVIDIA GPU drivers are currently installed:

```
nvidia-smi
```

This command should return the message nvidia-smi: not found.

- c. Run the following commands to enable maintenance mode on the host and install the NVIDIA vGPU Manager from the VIB file:

```
esxcli system maintenanceMode set --enable true  
esxcli software vib install -v /NVIDIA**.vib
```

You should see the message Operation finished successfully.

- d. Run the following command and verify that all eight GPU drivers are listed in the command output:

```
nvidia-smi
```

- e. Run the following command to verify that the NVIDIA vGPU package was installed and loaded correctly:

```
vmkload_mod -l | grep nvidia
```

The command should return output similar to the following: nvidia 816 13808

- f. Run the following command to reboot the host:

```
reboot -f
```

- g. Run the following command to exit maintenance mode:

```
esxcli system maintenanceMode set --enable false
```

6. Repeat steps 4-6 for any other newly deployed compute nodes with NVIDIA GPUs.
7. Perform the following tasks using the instructions in the NVIDIA documentation site:
 - a. Install the NVIDIA license server.
 - b. Configure the virtual machine guests for NVIDIA vGPU software.
 - c. If you are using vGPU-enabled desktops in a virtual desktop infrastructure (VDI) context, configure VMware Horizon View for NVIDIA vGPU software.

Find more information

- [NetApp HCI Resources page](#)

- [SolidFire and Element Software Documentation Center](#)

Configure Fully Qualified Domain Name web UI access

NetApp HCI with Element 12.2 or later enables you to access storage cluster web interfaces using the Fully Qualified Domain Name (FQDN). If you want to use the FQDN to access web user interfaces such as the Element web UI, per-node UI, or management node UI, you must first add a storage cluster setting to identify the FQDN used by the cluster. This enables the cluster to properly redirect a login session and improves integration with external services such as key managers and identity providers for multi-factor authentication.

What you'll need

- This feature requires Element 12.2 or later.
- Configuring this feature using NetApp Hybrid Cloud Control REST APIs requires management services 2.15 or later.
- Configuring this feature using the NetApp Hybrid Cloud Control UI requires management services 2.19 or later.
- To use REST APIs, you must have deployed a management node running version 11.5 or later.
- You need fully qualified domain names for the management node and each storage cluster that resolve correctly to the management node IP address and each storage cluster IP address.

You can configure or remove FQDN web UI access using NetApp Hybrid Cloud Control and the REST API. You can also troubleshoot incorrectly configured FQDNs.

- [Configure FQDN web UI access using NetApp Hybrid Cloud Control](#)
- [Configure FQDN web UI access using the REST API](#)
- [Remove FQDN web UI access using NetApp Hybrid Cloud Control](#)
- [Remove FQDN web UI access using the REST API](#)
- [Troubleshooting](#)

Configure FQDN web UI access using NetApp Hybrid Cloud Control

Steps

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select the menu icon at the top right of the page.
4. Select **Configure**.
5. In the **Fully Qualified Domain Names** pane, select **Set Up**.
6. In the resulting window, enter the FQDNs for the management node and each storage cluster.
7. Select **Save**.

The **Fully Qualified Domain Names** pane lists each storage cluster with its associated MVIP and FQDN.



Only connected storage clusters with the FQDN set are listed in the **Fully Qualified Domain Names** pane.

Configure FQDN web UI access using the REST API

Steps

1. Ensure that the Element storage nodes and the management node have DNS configured correctly for the network environment so that FQDNs in the environment can be resolved. To set DNS, go to the per-node UI for storage nodes and to the management node, then select **Network Settings > Management Network**.

- Per-node UI for storage nodes: https://<storage_node_management_IP>:442
- Per-node UI for the management node: <https://<ManagementNodeIP>:442>

2. Change the storage cluster settings using the Element API.

- Access the Element API and create the following cluster interface preference using the `CreateClusterInterfacePreference` API method, and insert the cluster MVIP FQDN for the preference value:
 - Name: `mvip_fqdn`
 - Value: <Fully Qualified Domain Name for the Cluster MVIP>

For example, the FQDN here is `storagecluster.my.org`:

```
https://<Cluster_MVIP>/json-rpc/12.2?method=CreateClusterInterfacePreference&name=mvip_fqdn&value=storagecluster.my.org
```

3. Change the management node settings using the REST API on the management node:

- Access the REST API UI for the management node by entering the management node IP address followed by `/mnode/2/`. For example:

```
https://<ManagementNodeIP>/mnode/2/
```

- Select **Authorize** or any lock icon and enter the Element cluster user name and password.
- Enter the client ID as `mnode-client`.
- Select **Authorize** to begin a session.
- Close the window.
- Select **GET /settings**.
- Select **Try it out**.
- Select **Execute**.
- Note whether or not the proxy is used as indicated in `"use_proxy"` by `true` or `false`.

- j. Select **PUT /settings**.
- k. Select **Try it out**.
- l. In the request body area, enter the management node FQDN as the value for the `mnode_fqdn` parameter. Also specify whether the proxy should be used (`true` or `false` from the previous step) for the `use_proxy` parameter.

```
{
  "mnode_fqdn": "mnode.my.org",
  "use_proxy": false
}
```

- m. Select **Execute**.

Remove FQDN web UI access using NetApp Hybrid Cloud Control

You can use this procedure to remove FQDN web access for the management node and the storage clusters.

Steps

1. In the **Fully Qualified Domain Names** pane, select **Edit**.
2. In the resulting window, delete the contents in the **FQDN** text field.
3. Select **Save**.

The window closes and the FQDN is no longer listed in the **Fully Qualified Domain Names** pane.

Remove FQDN web UI access using the REST API

Steps

1. Change the storage cluster settings using the Element API.
 - a. Access the Element API and delete the following cluster interface preference using the `DeleteClusterInterfacePreference` API method:
 - Name: `mvip_fqdn`

For example:

```
https://<Cluster_MVIP>/json-
rpc/12.2?method>DeleteClusterInterfacePreference&name=mvip_fqdn
```

2. Change the management node settings using the REST API on the management node:
 - a. Access the REST API UI for the management node by entering the management node IP address followed by `/mnode/2/`. For example:

```
https://<ManagementNodeIP>/mnode/2/
```

- b. Select **Authorize** or any lock icon and enter the Element cluster user name and password.
- c. Enter the client ID as `mnode-client`.
- d. Select **Authorize** to begin a session.
- e. Close the window.
- f. Select **PUT /settings**.
- g. Select **Try it out**.
- h. In the request body area, do not enter a value for the `mnode_fqdn` parameter. Also specify whether the proxy should be used (`true` or `false`) for the `use_proxy` parameter.

```
{
  "mnode_fqdn": "",
  "use_proxy": false
}
```

- i. Select **Execute**.

Troubleshooting

If FQDNs are configured incorrectly, you might have problems accessing either the management node, a storage cluster, or both. Use the following information to help troubleshoot the issue.

Issue	Cause	Resolution
<ul style="list-style-type: none"> You get a browser error when attempting to access either the management node or the storage cluster using the FQDN. You cannot log in to either the management node or the storage cluster using an IP address. 	The management node FQDN and storage cluster FQDN are both incorrectly configured.	Use the REST API instructions on this page to remove the management node and storage cluster FQDN settings and configure them again.
<ul style="list-style-type: none"> You get a browser error when attempting to access the storage cluster FQDN. You cannot log in to either the management node or the storage cluster using an IP address. 	The management node FQDN is correctly configured, but the storage cluster FQDN is incorrectly configured.	Use the REST API instructions on this page to remove the storage cluster FQDN settings and configure them again.

Issue	Cause	Resolution
<ul style="list-style-type: none"> • You get a browser error when attempting to access the management node FQDN. • You can log in to the management node and storage cluster using an IP address. 	The management node FQDN is incorrectly configured, but the storage cluster FQDN is correctly configured.	Log in to NetApp Hybrid Cloud Control to correct the management node FQDN settings in the UI, or use the REST API instructions on this page to correct the settings.

Find more information

- [CreateClusterInterfacePreference API](#) information in the [SolidFire and Element Documentation Center](#)
- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Access NetApp Hybrid Cloud Control

NetApp Hybrid Cloud Control enables you to manage NetApp HCI. You can upgrade management services and other components of NetApp HCI and expand and monitor your installation. You log in to NetApp Hybrid Cloud Control by browsing to the IP address of the management node.

What you'll need

- **Cluster administrator permissions:** You have permissions as administrator on the storage cluster.
- **Management services:** You have upgraded your management services to at least version 2.1.326. NetApp Hybrid Cloud Control is not available in earlier service bundle versions. For information about the current service bundle version, see the [Management Services Release Notes](#).

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.

The NetApp Hybrid Cloud Control interface appears.



If you logged in using insufficient permissions, you will see an "Unable to load" message throughout HCC resource pages and resources will not be available.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Reduce boot media wear on a NetApp HCI compute node

When you use flash memory or NVDIMM boot media with a NetApp HCI compute node, keeping the system logs on that media results in frequent writes to that media. This can eventually degrade the flash memory. Use the instructions in the following KB article to move host logging and the core dump file to a shared storage location, which can help prevent degradation of the boot media over time and help prevent full boot disk errors.

[How to reduce wear on the boot drive of a NetApp HCI compute node](#)

Find more information

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

Manage NetApp HCI

NetApp HCI management overview

You can manage credentials for NetApp HCI, user accounts, storage clusters, volumes, volume access groups, initiators, volume QoS policies and the management node.

Here are the items you can work with:

- [Update vCenter and ESXi credentials](#)
- [Manage NetApp HCI storage assets](#)
- [Work with the management node](#)
- [Power your NetApp HCI system off or on](#)

Update vCenter and ESXi credentials

To maintain full functionality of NetApp Hybrid Cloud Control for your NetApp HCI installation, when you change your credentials in vCenter and ESXi hosts, you also need to update those credentials in the asset service on the management node.

About this task

NetApp Hybrid Cloud Control communicates with vCenter and the individual compute nodes running VMware vSphere ESXi to retrieve information for the dashboard and to facilitate rolling upgrades of firmware, software and drivers. NetApp Hybrid Cloud Control and its related services on the management node use credentials (username/password) to authenticate against VMware vCenter and ESXi.

If communication between these components fails, NetApp Hybrid Cloud Control and vCenter display error messages when authentication problems occur. NetApp Hybrid Cloud Control will display a red error banner if it cannot communicate with the associated VMware vCenter instance in the NetApp HCI installation. VMware vCenter will display ESXi account lockout messages for individual ESXi hosts as a result of NetApp Hybrid Cloud Control using outdated credentials.

The management node in NetApp HCI refers to these components using the following names:

- "Controller assets" are vCenter instances associated with your NetApp HCI installation.
- "Compute node assets" are the ESXi hosts in your NetApp HCI installation.

During the initial installation of NetApp HCI using the NetApp Deployment Engine, the management node stored the credentials for the administrative user you specified for vCenter and the "root" account password on ESXi servers.

Update vCenter password by using the management node REST API

Follow the steps to update the controller assets. See [View or edit existing controller assets](#).

Update the ESXi password by using the management node REST API

Steps

1. To gain an overview of the Management node REST API user interface, see the [Management node REST API user interface overview](#).
2. Access the REST API UI for management services on the management node:

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

Replace <management node IP> with the IPv4 address of your management node on the management network used for NetApp HCI.

3. Select **Authorize** or any lock icon and complete the following:
 - a. Enter the NetApp SolidFire cluster administrative user name and password.
 - b. Enter the client ID as `mnode-client`.
 - c. Select **Authorize** to begin a session.
 - d. Close the window.
4. From the REST API UI, select **GET /assets/compute_nodes**.

This retrieves the records of compute node assets that are stored in the management node.

Here is the direct link to this API in the UI:

```
https://<ManagementNodeIP>/mnode/#/assets/routes.v1.assets_api.get_compu  
te_nodes
```

5. Select **Try it out**.
6. Select **Execute**.
7. From the response body, identify the compute node asset records that need updated credentials. You can use the “ip” and “host_name” properties to find the correct ESXi host records.

```
"config": { },  
"credentialid": <credential_id>,  
"hardware_tag": <tag>,  
"host_name": <host_name>,  
"id": <id>,  
"ip": <ip>,  
"parent": <parent>,  
"type": ESXi Host
```



The next step uses the “parent” and “id” fields in the compute asset record to reference the record to be updated.

8. Configure the specific compute node asset:
 - a. Select **PUT /assets/{asset_id}/compute-nodes/{compute_id}**.

Here is the direct link to the API in the UI:

```
https://<ManagementNodeIP>/mnode/#/assets/routes.v1.assets_api.put_assets_compute_id
```

- b. Select **Try it out**.
- c. Enter the “asset_id” with the “parent” information.
- d. Enter the “compute_id” with the “id” information.
- e. Modify the request body in the user interface to update only the password and user name parameters in the compute asset record:

```
{  
  "password": "<password>",  
  "username": "<username>"  
}
```

- f. Select **Execute**.
- g. Validate that the response is HTTP 200, which indicates that the new credentials have been stored in the referenced compute asset record
9. Repeat the previous two steps for additional compute node assets that need to be updated with a new password.
10. Navigate to https://<mNode_ip>/inventory/1/.
 - a. Select **Authorize** or any lock icon and complete the following:
 - i. Enter the NetApp SolidFire cluster administrative user name and password.
 - ii. Enter the client ID as mnode-client.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the window.
 - b. From the REST API UI, select **GET /installations**.
 - c. Select **Try it out**.
 - d. Select **True** from the refresh description drop-down list.
 - e. Select **Execute**.
 - f. Validate that the response is HTTP 200.
11. Wait for about 15 minutes for the account lockout message in vCenter to disappear.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Manage NetApp HCI storage

Manage NetApp HCI storage overview

With NetApp HCI, you can manage these storage assets by using the NetApp Hybrid Cloud Control.

- [Create and manage user accounts](#)
- [Add and manage storage clusters](#)
- [Create and manage volumes](#)
- [Create and manage volume access groups](#)
- [Create and manage initiators](#)
- [Create and manage volume QoS policies](#)

Find more information

- [SolidFire and Element 12.2 Documentation Center](#)
- [NetApp Element Plug-in for vCenter Server](#)

Create and manage user accounts by using NetApp Hybrid Cloud Control

In Element-based storage systems, authoritative cluster users can be created to enable login access to NetApp Hybrid Cloud Control depending on the permissions you want to grant "Administrator" or "Read-only" users. In addition to cluster users, there are also volume accounts, which enable clients to connect to volumes on a storage node.

Manage the following types of accounts:

- [Manage authoritative cluster accounts](#)
- [Manage volume accounts](#)

Enable LDAP

To use LDAP for any user account, you must first enable LDAP.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select on the top right Options icon and select **User Management**.
3. From the Users page, select **Configure LDAP**.
4. Define your LDAP configuration.
5. Select the authentication type of Search and Bind or Direct Bind.
6. Before you save the changes, select **Test LDAP Log In** at the top of the page, enter the user name and password of a user you know exists, and select **Test**.
7. Select **Save**.

Manage authoritative cluster accounts

[Authoritative user accounts](#) are managed from the top right menu User Management option in NetApp Hybrid Cloud Control. These types of accounts enable you to authenticate against any storage asset associated with a NetApp Hybrid Cloud Control instance of nodes and clusters. With this account, you can manage volumes, accounts, access groups, and more across all clusters.

Create an authoritative cluster account

You can create an account by using NetApp Hybrid Cloud Control.

This account can be used to log in to the Hybrid Cloud Control, the per-node UI for the cluster, and the storage cluster in NetApp Element software.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select the top right Options icon, then select **User Management**.
3. Select **Create User**.
4. Select the authentication type of cluster or LDAP.
5. Complete one of the following:
 - If you selected LDAP, enter the DN.



To use LDAP, you must first enable LDAP or LDAPS. See [Enable LDAP](#).

- If you selected Cluster as the Auth Type, enter a name and password for the new account.

6. Select either Administrator or Read-only permissions.



To view the permissions from NetApp Element software, select **Show legacy permissions**. If you select a subset of these permissions, the account is assigned Read-only permissions. If you select all legacy permissions, the account is assigned Administrator permissions.



To ensure that all children of a group inherit permissions, create a DN organization admin group in the LDAP server. All the children accounts of that group will inherit those permissions.

7. Check the box indicating that "I have read and accept the NetApp End User License Agreement."
8. Select **Create User**.

Edit an authoritative cluster account

You can change the permissions or password on a user account by using NetApp Hybrid Cloud Control.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select on the icon in the top right and select **User Management**.
3. Optionally filter the list of user accounts by selecting **Cluster**, **LDAP**, or **Idp**.

If you configured users on the storage cluster with LDAP, those accounts show a User Type of "LDAP." If you configured users on the storage cluster with Idp, those accounts show a User Type of "Idp."

4. In the **Actions** column in the table, expand the menu for the account and select **Edit**.
5. Make changes as needed.
6. Select **Save**.
7. Log out of NetApp Hybrid Cloud Control.
8. [Update the credentials](#) for the authoritative cluster asset using the NetApp Hybrid Cloud Control API.



It might take the NetApp Hybrid Cloud Control UI up to 15 minutes to refresh the inventory. To manually refresh inventory, access the REST API UI inventory service <https://<ManagementNodeIP>/inventory/1/> and run GET /installations/{id} for the cluster.

9. Log into NetApp Hybrid Cloud Control.

Delete an authoritative user account

You can delete one or more accounts when it is no longer needed. You can delete an LDAP user account.

You cannot delete the primary administrator user account for the authoritative cluster.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select on the icon in the top right and select **User Management**.
3. In the **Actions** column in the Users table, expand the menu for the account and select **Delete**.
4. Confirm the deletion by selecting **Yes**.

Manage volume accounts

[Volume accounts](#) are managed within the NetApp Hybrid Cloud Control Volumes table. These accounts are specific only to the storage cluster on which they were created. These types of accounts enable you to set permissions on volumes across the network, but have no effect outside of those volumes.

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

Create a volume account

Create an account specific to this volume.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select **Storage > Volumes**.
3. Select the **Accounts** tab.
4. Select the **Create Account** button.
5. Enter a name for the new account.

6. In the CHAP Settings section, enter the following information:

- Initiator Secret for CHAP node session authentication
- Target Secret for CHAP node session authentication



To auto-generate either password, leave the credential fields blank.

7. Select **Create Account**.

Edit a volume account

You can change the CHAP info and change whether an account is active or locked.



Deleting or locking an account associated with the management node results in an inaccessible management node.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select **Storage > Volumes**.
3. Select the **Accounts** tab.
4. In the **Actions** column in the table, expand the menu for the account and select **Edit**.
5. Make changes as needed.
6. Confirm the changes by selecting **Yes**.

Delete a volume account

Delete an account that you no longer need.

Before you delete a volume account, delete and purge any volumes associated with the account first.



Deleting or locking an account associated with the management node results in an inaccessible management node.



Persistent volumes that are associated with management services are assigned to a new account during installation or upgrade. If you are using persistent volumes, do not modify or delete the volumes or their associated account. If you do delete these accounts, you could render your management node unusable.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select **Storage > Volumes**.
3. Select the **Accounts** tab.
4. In the **Actions** column in the table, expand the menu for the account and select **Delete**.
5. Confirm the deletion by selecting **Yes**.

Find more information

- [Learn about accounts](#)
- [Work with user accounts](#)
- [NetApp Element Plug-in for vCenter Server](#)

Add and manage storage clusters using NetApp Hybrid Cloud Control

You can add storage clusters to the management node assets inventory so that they can be managed using NetApp Hybrid Cloud Control (HCC). The first storage cluster added during system setup is the default [authoritative storage cluster](#), but additional clusters can be added using HCC UI.

After a storage cluster is added, you can monitor cluster performance, change storage cluster credentials for the managed asset, or remove a storage cluster from the management node asset inventory if it no longer needs to be managed using HCC.

Starting with Element 12.2, you can use the [maintenance mode](#) feature options to enable and disable maintenance mode for your storage cluster nodes.

What you'll need

- **Cluster administrator permissions:** You have permissions as administrator on the [authoritative storage cluster](#). The authoritative cluster is the first cluster added to the management node inventory during system setup.
- **Element software:** Your storage cluster version is running NetApp Element software 11.3 or later.
- **Management node:** You have deployed a management node running version 11.3 or later.
- **Management services:** You have updated your management services bundle to version 2.17 or later.

Options

- [Add a storage cluster](#)
- [Confirm storage cluster status](#)
- [Edit storage cluster credentials](#)
- [Remove a storage cluster](#)
- [Enable and disable maintenance mode](#)

Add a storage cluster

You can add a storage cluster to the management node assets inventory using NetApp Hybrid Cloud Control. This allows you to manage and monitor the cluster using the HCC UI.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the authoritative storage cluster administrator credentials.
2. From the Dashboard, select the options menu on the top right and select **Configure**.
3. In the **Storage Clusters** pane, select **Storage Cluster Details**.
4. Select **Add Storage Cluster**.
5. Enter the following information:
 - Storage cluster management virtual IP address



Only remote storage clusters that are not currently managed by a management node can be added.

- Storage cluster user name and password

6. Select **Add**.



After you add the storage cluster, the cluster inventory can take up to 15 minutes to refresh and display the new addition. You might need to refresh the page in your browser to see the changes.

7. If you are adding Element eSDS clusters, enter or upload your SSH private key and SSH user account.

Confirm storage cluster status

You can monitor the connection status of storage clusters assets using the NetApp Hybrid Cloud Control UI.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the authoritative storage cluster administrator credentials.
2. From the Dashboard, select the options menu on the top right and select **Configure**.
3. Review the status of storage clusters in the inventory.
4. From the **Storage Clusters** pane, select **Storage Cluster Details** for additional detail.

Edit storage cluster credentials

You can edit the storage cluster's administrator user name and password using the NetApp Hybrid Cloud Control UI.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the authoritative storage cluster administrator credentials.
2. From the Dashboard, select the options menu on the top right and select **Configure**.
3. In the **Storage Clusters** pane, select **Storage Cluster Details**.
4. Select the **Actions** menu for the cluster and select **Edit Cluster Credentials**.
5. Update the storage cluster user name and password.
6. Select **Save**.

Remove a storage cluster

Removing a storage cluster from NetApp Hybrid Cloud Control removes the cluster from the management node inventory. After you remove a storage cluster, the cluster can no longer be managed by HCC and you can access it only by navigating directly to its management IP address.



You cannot remove the authoritative cluster from the inventory. To determine the authoritative cluster, go to **User Management > Users**. The authoritative cluster is listed next to the heading **Users**.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the authoritative storage cluster administrator credentials.
2. From the Dashboard, select the options menu on the top right and select **Configure**.
3. In the **Storage Clusters** pane, select **Storage Cluster Details**.
4. Select the **Actions** menu for the cluster and select **Remove Storage Cluster**.



Selecting **Yes** next removes the cluster from the installation.

5. Select **Yes**.

Enable and disable maintenance mode

This [maintenance mode](#) feature options give you the capability to [enable](#) and [disable](#) maintenance mode for a storage cluster node.

What you'll need

- **Element software:** Your storage cluster version is running NetApp Element software 12.2 or later.
- **Management node:** You have deployed a management node running version 12.2 or later.
- **Management services:** You have updated your management services bundle to version 2.19 or later.
- You have access to log in at the administrator level.

Enable maintenance mode

You can use the following procedure to enable maintenance mode for a storage cluster node.



Only one node can be in maintenance mode at a time.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

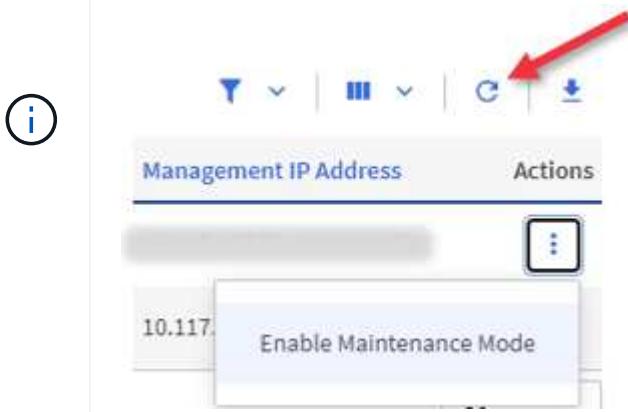
2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.



The maintenance mode feature options are disabled at the read-only level.

3. In the left navigation blue box, select the NetApp HCI installation.
4. In the left navigation pane, select **Nodes**.
5. To view storage inventory information, select **Storage**.
6. Enable maintenance mode on a storage node:

The storage nodes table is updated automatically every two minutes for non-user initiated actions. Before an action, to ensure that you have the most up-to-date status, you can refresh the nodes table by using the refresh icon located on the upper-right side of the nodes table.



- a. Under **Actions**, select **Enable Maintenance Mode**.

While **Maintenance Mode** is being enabled, maintenance mode actions are unavailable for the selected node and all other nodes on the same cluster.

After **Enabling Maintenance Mode** completes, the **Node Status** column displays a wrench icon and the text "**Maintenance Mode**" for the node that is in maintenance mode.

Disable maintenance mode

After a node is successfully placed in maintenance mode, the **Disable Maintenance Mode** action is available for this node. Actions on the other nodes are unavailable until maintenance mode is disabled successfully on the node undergoing maintenance.

Steps

1. For the node under maintenance mode, under **Actions**, select **Disable Maintenance Mode**.

While **Maintenance Mode** is being disabled, maintenance mode actions are unavailable for the selected node and all other nodes on the same cluster.

After **Disabling Maintenance Mode** completes, the **Node Status** column displays **Active**.



When a node is in maintenance mode, it does not accept new data. As a result, it can take longer to disable maintenance mode because the node must sync its data back up before it can exit maintenance mode. The longer you spend in maintenance mode, the longer it can take to disable maintenance mode.

Troubleshoot

If you encounter errors when you are either enabling or disabling maintenance mode, a banner error displays at the top of the nodes table. For more information on the error, you can select the **Show Details** link that is provided on the banner to show what the API returns are.

Find more information

- [Create and manage storage cluster assets](#)

Create and manage volumes by using NetApp Hybrid Cloud Control

You can create a volume and associate the volume with a given account. Associating a volume with an account gives the account access to the volume through the iSCSI initiators and CHAP credentials.

You can specify QoS settings for a volume during creation.

You can manage volumes in NetApp Hybrid Cloud Control in the following ways:

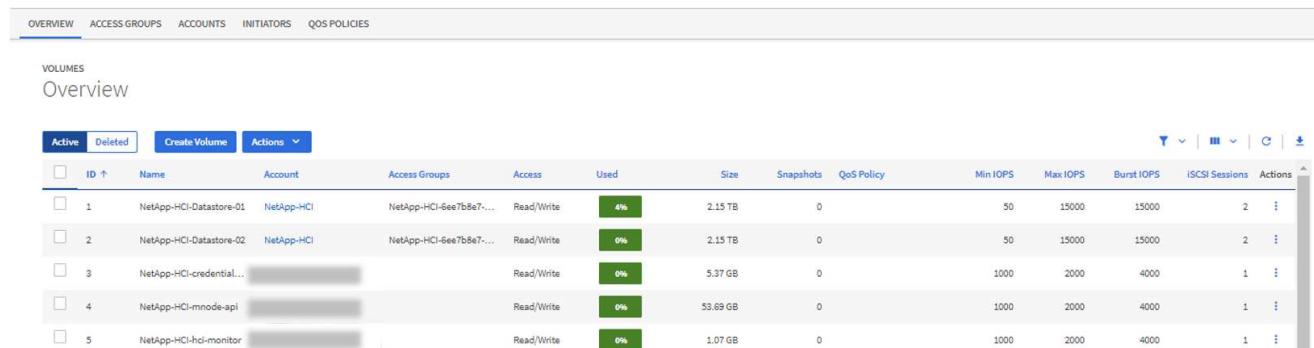
- [Create a volume](#)
- [Apply a QoS policy to a volume](#)
- [Edit a volume](#)
- [Clone volumes](#)
- [Delete a volume](#)
- [Restore a deleted volume](#)
- [Purge a deleted volume](#)

Create a volume

You can create a storage volume using NetApp Hybrid Cloud Control.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes > Overview** tab.



Active	Deleted	Create Volume	Actions										
ID	Name	Account	Access Groups	Access	Used	Size	Snapshots	QoS Policy	Min IOPS	Max IOPS	Burst IOPS	iSCSI Sessions	Actions
<input type="checkbox"/>	1	NetApp-HCI-Datastore-01	NetApp-HCI	NetApp-HCI-6ee7b8e7...	Read/Write	<div style="width: 4%;">4%</div>	2.15 TB	0	50	15000	15000	2	⋮
<input type="checkbox"/>	2	NetApp-HCI-Datastore-02	NetApp-HCI	NetApp-HCI-6ee7b8e7...	Read/Write	<div style="width: 0%;">0%</div>	2.15 TB	0	50	15000	15000	2	⋮
<input type="checkbox"/>	3	NetApp-HCI-credential...			Read/Write	<div style="width: 0%;">0%</div>	5.37 GB	0	1000	2000	4000	1	⋮
<input type="checkbox"/>	4	NetApp-HCI-mnode-api			Read/Write	<div style="width: 0%;">0%</div>	53.69 GB	0	1000	2000	4000	1	⋮
<input type="checkbox"/>	5	NetApp-HCI-hci-monitor			Read/Write	<div style="width: 0%;">0%</div>	1.07 GB	0	1000	2000	4000	1	⋮

4. Select **Create Volume**.
5. Enter a name for the new volume.
6. Enter the total size of the volume.



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

7. Select a block size for the volume.
8. From the Account list, select the account that should have access to the volume.

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



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

9. To set the Quality of Service, do one of the following:
 - a. Select an existing QoS policy.
 - b. Under QoS Settings, set customized minimum, maximum, and burst values for IOPS or use the default QoS values.

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

10. Select **Create Volume**.

Apply a QoS policy to a volume

You can apply a QoS policy to an existing storage volume by using NetApp Hybrid Cloud Control.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes > Overview**.
4. In the **Actions** column in the volumes table, expand the menu for the volume and select **Edit**.
5. Change the Quality of Service by doing one of the following:
 - a. Select an existing policy.
 - b. Under Custom Settings, set the minimum, maximum, and burst values for IOPS or use the default values.



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



When you change IOPS values, increment in tens or hundreds. Input values require valid whole numbers. Configure volumes with an extremely high burst value. This enables the system to process occasional large block, sequential workloads more quickly, while still constraining the sustained IOPS for a volume.

6. Select **Save**.

Edit a volume

Using NetApp Hybrid Cloud Control, you can edit volume attributes such as QoS values, volume size, and the unit of measurement by which byte values are calculated. You can also modify account access for replication usage or to restrict access to the volume.

About this task

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

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

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes > Overview**.
4. In the **Actions** column in the volumes table, expand the menu for the volume and select **Edit**.
5. Make changes as needed:

- a. Change the total size of the volume.



You can increase, but not decrease, the size of the volume. You can only resize one volume in a single resizing operation. Garbage collection operations and software upgrades do not interrupt the resizing operation.



If you are adjusting volume size for replication, first increase the size of the volume assigned as the replication target. Then you can resize the source volume. The target volume can be greater or equal in size to the source volume, but it cannot be smaller.



The default volume size selection is in GB. You can create volumes using sizes measured in GB or GiB:
1GB = 1 000 000 000 bytes
1GiB = 1 073 741 824 bytes

- b. Select a different account access level:

- Read Only
- Read/Write
- Locked
- Replication Target

- c. Select the account that should have access to the volume.

Begin typing and the auto-complete function displays possible values for you to choose.

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

d. Change the Quality of Service by doing one of the following:

- i. Select an existing policy.
- ii. Under Custom Settings, set the minimum, maximum, and burst values for IOPS or use the default values.



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



When you change IOPS values, you should increment in tens or hundreds. Input values require valid whole numbers. Configure volumes with an extremely high burst value. This enables the system to process occasional large block, sequential workloads more quickly, while still constraining the sustained IOPS for a volume.

6. Select **Save**.

Clone volumes

You can create a clone of a single storage volume or clone a group of volumes to make a point-in-time copy of the data. When you clone a volume, the system creates a snapshot of the volume and then creates a copy of the data referenced by the snapshot.

What you'll need

- At least one cluster must be added and running.
- At least one volume has been 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.

Volume cloning 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.



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

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select the **Volumes > Overview** tab.
4. Select each volume you want to clone and select the **Clone** button that appears.
5. Do one of the following:
 - To clone a single volume, perform the following steps:

- a. In the **Clone Volume** dialog box, enter a volume name for the volume clone.



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

- b. Select an account access level:

- Read Only
- Read/Write
- Locked
- Replication Target

- c. Select a size in GB or GIB for the volume clone.



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.

- d. Select an account to associate with the volume clone.

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

- e. Select **Clone Volumes**.

- To clone multiple volumes, perform the following steps:

- a. In the **Clone Volumes** dialog box, enter an optional prefix for the volume clones in the **New Volume Name Prefix** field.
- b. Select a new type of access for the volume clones or copy the access type from the active volumes.
- c. Select a new account to associate with the volume clones or copy the account association from the active volumes.
- d. Select **Clone Volumes**.



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.

Delete a volume

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

About this task

The system does not immediately purge deleted volumes; they remain available for approximately eight hours. After eight hours, they are purged and no longer available. If you restore a volume before the system purges it, the volume comes back online and iSCSI connections are restored.

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



Persistent volumes that are associated with management services are created and assigned to a new account during installation or upgrade. If you are using persistent volumes, do not modify or delete the volumes or their associated account. If you do delete these volumes, you could render your management node unusable.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes > Overview**.
4. Select one or more volumes to delete.
5. Do one of the following:
 - If you selected multiple volumes, select the **Delete** quick filter at the top of the table.
 - If you selected a single volume, in the **Actions** column of the Volumes table, expand the menu for the volume and select **Delete**.
6. Confirm the delete by selecting **Yes**.

Restore a deleted volume

After a storage volume is deleted, you can still restore it if you do so before eight hours after deletion.

The system does not immediately purge deleted volumes; they remain available for approximately eight hours. After eight hours, they are purged and no longer available. If you restore a volume before the system purges it, the volume comes back online and iSCSI connections are restored.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes > Overview**.
4. Select **Deleted**.
5. In the **Actions** column of the Volumes table, expand the menu for the volume and select **Restore**.
6. Confirm the process by selecting **Yes**.

Purge a deleted volume

After storage volumes are deleted, they remain available for approximately eight hours. After eight hours, they are purged automatically and no longer available. If you do not want to wait for the eight hours, you can delete

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes > Overview**.
4. Select **Deleted**.

5. Select one or more volumes to purge.
6. Do one of the following:
 - If you selected multiple volumes, select the **Purge** quick filter at the top of the table.
 - If you selected a single volume, in the **Actions** column of the Volumes table, expand the menu for the volume and select **Purge**.
7. In the **Actions** column of the Volumes table, expand the menu for the volume and select **Purge**.
8. Confirm the process by selecting **Yes**.

Find more information

- [Learn about volumes](#)
- [Work with volumes](#)
- [NetApp Element Plug-in for vCenter Server](#)

Create and manage volume access groups

You can create new volume access groups, make changes to the name, associated initiators, or associated volumes of access groups, or delete existing volume access groups using NetApp Hybrid Cloud Control.

What you'll need

- You have administrator credentials for this NetApp HCI system.
- You have upgraded your management services to at least version 2.15.28. NetApp Hybrid Cloud Control storage management is not available in earlier service bundle versions.
- Ensure you have a logical naming scheme for volume access groups.

Add a volume access group

You can add a volume access group to a storage cluster by using NetApp Hybrid Cloud Control.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes**.
4. Select the **Access Groups** tab.
5. Select the **Create Access Group** button.
6. In the resulting dialog, enter a name for the new volume access group.
7. (Optional) In the **Initiators** section, select one or more initiators to associate with the new volume access group.

If you associate an initiator with the volume access group, that initiator can access each volume in the group without the need for authentication.

8. (Optional) In the **Volumes** section, select one or more volumes to include in this volume access group.

9. Select **Create Access Group**.

Edit a volume access group

You can edit the properties of an existing volume access group by using NetApp Hybrid Cloud Control. You can make changes to the name, associated initiators, or associated volumes of an access group.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes**.
4. Select the **Access Groups** tab.
5. In the **Actions** column of the table of access groups, expand the options menu for the access group you need to edit.
6. In the options menu, select **Edit**.
7. Make any needed changes to the name, associated initiators, or associated volumes.
8. Confirm your changes by selecting **Save**.
9. In the **Access Groups** table, verify that the access group reflects your changes.

Delete a volume access group

You can remove a volume access group by using NetApp Hybrid Cloud Control, and at the same time remove the initiators associated with this access group from the system.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes**.
4. Select the **Access Groups** tab.
5. In the **Actions** column of the table of access groups, expand the options menu for the access group you need to delete.
6. In the options menu, select **Delete**.
7. If you do not wish to delete the initiators that are associated with the access group, deselect the **Delete initiators in this access group** checkbox.
8. Confirm the delete operation by selecting **Yes**.

Find more information

- [Learn about volume access groups](#)
- [Add initiator to a volume access group](#)
- [NetApp Element Plug-in for vCenter Server](#)

Create and manage initiators

You can use [initiators](#) for CHAP-based rather than account-based access to 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.

What you'll need

- You have cluster administrator credentials.
- You have upgraded your management services to at least version 2.17. NetApp Hybrid Cloud Control initiator management is not available in earlier service bundle versions.

Options

- [Create an initiator](#)
- [Add initiators to a volume access group](#)
- [Change an initiator alias](#)
- [Delete initiators](#)

Create an initiator

You can create iSCSI or Fibre Channel initiators and optionally assign them aliases.

About this task

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
```

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes**.
4. Select the **Initiators** tab.
5. Select the **Create Initiators** button.

Option	Steps
Create one or more initiators	<ol style="list-style-type: none"> Enter the IQN or WWPN for the initiator in the IQN/WWPN field. Enter a friendly name for the initiator in the Alias field. (Optional) Select Add Initiator to open new initiator fields or use the bulk create option instead. Select Create Initiators.
Bulk create initiators	<ol style="list-style-type: none"> Select Bulk Add IQNs/WWPNs. Enter a list of IQNs or WWPNs in the text box. Each IQN or WWPN must be comma or space separated or on its own line. Select Add IQNs/WWPNs. (Optional) Add unique aliases to each initiator. Remove any initiator from the list that might already exist in the installation. Select Create Initiators.

Add initiators to a volume access group

You can add initiators to an volume access group. When you add an initiator to a volume access group, the initiator enables access to all volumes in that volume access group.

Steps

- Log in to NetApp Hybrid Cloud Control by providing the Element storage cluster administrator credentials.
- From the Dashboard, expand the name of your storage cluster on the left navigation menu.
- Select **Volumes**.
- Select the **Initiators** tab.
- Select one or more initiators you want to add.
- Select **Actions > Add to Access Group**.
- Select the access group.
- Confirm your changes by selecting **Add Initiator**.

Change an initiator alias

You can change the alias of an existing initiator or add an alias if one does not already exist.

Steps

- Log in to NetApp Hybrid Cloud Control by providing the Element storage cluster administrator credentials.
- From the Dashboard, expand the name of your storage cluster on the left navigation menu.
- Select **Volumes**.

4. Select the **Initiators** tab.
5. In the **Actions** column, expand the options menu for the initiator.
6. Select **Edit**.
7. Make any needed changes to the alias or add a new alias.
8. Select **Save**.

Delete initiators

You can delete one or more initiators. 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. Log in to NetApp Hybrid Cloud Control by providing the Element storage cluster administrator credentials.
2. From the Dashboard, expand the name of your storage cluster on the left navigation menu.
3. Select **Volumes**.
4. Select the **Initiators** tab.
5. Delete one or more initiators:
 - a. Select one or more initiators you want to delete.
 - b. Select **Actions > Delete**.
 - c. Confirm the delete operation and select **Yes**.

Find more information

- [Learn about initiators](#)
- [Learn about volume access groups](#)
- [NetApp Element Plug-in for vCenter Server](#)

Create and manage volume 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. The selected cluster must be Element 10.0 or later to use QoS policies; otherwise, QoS policy functions are not available.



See NetApp HCI Concepts content for more information about using [QoS policies](#) instead of individual volume [QoS](#).

Using NetApp Hybrid Cloud Control, you can create and manage QoS policies by completing the following tasks:

- [Create a QoS policy](#)
- [Apply a QoS policy to a volume](#)
- [Change the QoS policy assignment of a volume](#)
- [Edit a QoS policy](#)

- [Delete a QoS policy](#)

Create a QoS policy

You can create QoS policies and apply them to volumes that should have equivalent performance.



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

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the menu for your storage cluster.
3. Select **Storage > Volumes**.
4. Select the **QoS Policies** tab.
5. Select **Create Policy**.
6. Enter the **Policy Name**.



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

7. Enter the minimum IOPS, maximum IOPS, and burst IOPS values.
8. Select **Create QoS Policy**.

A system ID is generated for the policy and the policy appears on the QoS Policies page with its assigned QoS values.

Apply a QoS policy to a volume

You can assign an existing QoS policy to a volume using NetApp Hybrid Cloud Control.

What you'll need

The QoS policy you want to assign has been [created](#).

About this task

This task describes how to assign a QoS policy to an individual volume by changing its settings. The latest version of NetApp Hybrid Cloud Control does not have a bulk assign option for more than one volume. Until the functionality to bulk assign is provided in a future release, you can use the Element web UI or vCenter Plug-in UI to bulk assign QoS policies.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the menu for your storage cluster.
3. Select **Storage > Volumes**.
4. Select the **Actions** menu next to the volume you intend to modify.
5. In the resulting menu, select **Edit**.

6. In the dialog box, enable **Assign QoS Policy** and select the QoS policy from the drop-down list to apply to the selected volume.



Assigning QoS will override any individual volume QoS values that have been previously applied.

7. Select **Save**.

The updated volume with the assigned QoS policy appears on the Overview page.

Change the QoS policy assignment of a volume

You can remove the assignment of a QoS policy from a volume or select a different QoS policy or custom QoS.

What you'll need

The volume you want to modify is [assigned](#) a QoS policy.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the menu for your storage cluster.
3. Select **Storage > Volumes**.
4. Select the **Actions** menu next to the volume you intend to modify.
5. In the resulting menu, select **Edit**.
6. In the dialog box, do one of the following:
 - Disable **Assign QoS Policy** and modify the **Min IOPS**, **Max IOPS**, and **Burst IOPS** values for individual volume QoS.



When QoS policies are disabled, the volume uses default QoS IOPS values unless otherwise modified.

- Select a different QoS policy from the drop-down list to apply to the selected volume.

7. Select **Save**.

The updated volume appears on the Overview page.

Edit a QoS policy

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

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the menu for your storage cluster.
3. Select **Storage > Volumes**.
4. Select the **QoS Policies** tab.

5. Select the **Actions** menu next to the QoS policy you intend to modify.
6. Select **Edit**.
7. In the **Edit QoS Policy** dialog box, change one or more of the following:
 - **Name**: The user-defined name for the QoS policy.
 - **Min IOPS**: The minimum number of IOPS guaranteed for the volume. Default = 50.
 - **Max IOPS**: The maximum number of IOPS allowed for the volume. Default = 15,000.
 - **Burst IOPS**: The maximum number of IOPS allowed over a short period of time for the volume. Default = 15,000.
8. Select **Save**.

The updated QoS policy appears on the QoS Policies page.



You can select the link in the **Active Volumes** column for a policy to show a filtered list of the volumes assigned to that policy.

Delete a QoS policy

You can delete a QoS policy if it is no longer needed. When you delete a QoS policy, all volumes assigned with the policy maintain the QoS values previously defined by the policy but as individual volume QoS. Any association with the deleted QoS policy is removed.

Steps

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, expand the menu for your storage cluster.
3. Select **Storage > Volumes**.
4. Select the **QoS Policies** tab.
5. Select the **Actions** menu next to the QoS policy you intend to modify.
6. Select **Delete**.
7. Confirm the action.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp SolidFire and Element Documentation Center \(Documentation Center Versions\)](#)

Work with the management node

Management node overview

You can use the management node (mNode) to use system services, manage cluster assets and settings, run system tests and utilities, configure Active IQ for system monitoring, and enable NetApp Support access for troubleshooting.

For clusters running Element software version 11.3 or later, you can work with the management node by using

one of two interfaces:

- With the management node UI ([https://\[mNode IP\]:442](https://[mNode IP]:442)), you can make changes to network and cluster settings, run system tests, or use system utilities.
- With the built-in REST API UI ([https://\[mNode IP\]/mnode](https://[mNode IP]/mnode)), you can run or understand APIs relating to the management node services, including proxy server configuration, service level updates, or asset management.

Install or recover a management node:

- [Install a management node](#)
- [Configure a storage Network Interface Controller \(NIC\)](#)
- [Recover a management node](#)

Access the management node:

- [Access the management node \(UI or REST API\)](#)

Change the default SSL certificate:

- [Change the management node default SSL certificate](#)

Perform tasks with the management node UI:

- [Management node UI overview](#)

Perform tasks with the management node REST APIs:

- [Management node REST API UI overview](#)

Disable or enable remote SSH functionality or start a remote support tunnel session with NetApp Support to help you troubleshoot:

- [Enable remote NetApp Support connections](#)
- [Manage SSH functionality on the management node](#)

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Install or recover a management node

Install 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 NetApp HCI administrators who are not using the NetApp Deployment Engine for management node installation.

What you'll need

- Your cluster version is running NetApp Element software 11.3 or later.

- Your installation uses IPv4. The management node 11.3 does not support IPv6.



If you need to IPv6 support, you can use the management node 11.1.

- You have permission to download software from the NetApp Support Site.
- You have identified the management node image type that is correct for your platform:

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

- (Management node 12.0 and 12.2 with proxy server) You have updated NetApp Hybrid Cloud Control to management services version 2.16 before configuring a proxy server.

About this task

The Element 12.2 management node is an optional upgrade. It is not required for existing deployments.

Prior to following this procedure, you should have an understanding of [persistent volumes](#) and whether or not you want to use them. Persistent volumes are optional but recommended for management node configuration data recovery in the event of a virtual machine (VM) loss.

Steps

1. [Download ISO or OVA and deploy the VM](#)
2. [Create the management node admin and configure the network](#)
3. [Configure time sync](#)
4. [Set up the management node](#)
5. [Configure controller assets](#)
6. [\(NetApp HCI only\) Configure compute node assets](#)

Download ISO or OVA and deploy the VM

1. Download the OVA or ISO for your installation from the [NetApp HCI](#) page on the NetApp Support Site:
 - a. Select **Download Latest Release** and accept the EULA.
 - b. 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 (for example, 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 VM from your hypervisor with the following configuration:

- Six virtual CPUs
- 24GB RAM
- Storage adapter type set to LSI Logic Parallel



The default for your management node might be LSI Logic SAS. In the **New Virtual Machine** window, verify the storage adapter configuration by selecting **Customize hardware > Virtual Hardware**. If required, change LSI Logic SAS to **LSI Logic Parallel**.

- 400GB virtual disk, thin provisioned
- One virtual network interface with internet access and access to the storage MVIP.
- 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.



Do not power on the VM prior to the step indicating to do so later in this procedure.

- a. Attach the ISO to the VM and boot to the .iso install image.



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

4. Power on the VM for the management node after the installation completes.

Create the management node admin and configure the network

1. Using the terminal user interface (TUI), create a management node admin user.



To move through the menu options, press the Up or Down arrow keys. To move through the buttons, press Tab. To move from the buttons to the fields, press Tab. To navigate between fields, press the Up or Down arrow keys.

2. If there is a Dynamic Host Configuration Protocol (DHCP) server on the network that assigns IPs with a maximum transmission unit (MTU) less than 1500 bytes, you must perform the following steps:

- a. Temporarily put the management node on a vSphere network without DHCP, such as iSCSI.
- b. Reboot the VM or restart the VM network.
- c. Using the TUI, configure the correct IP on the management network with an MTU greater than or equal to 1500 bytes.
- d. Re-assign the correct VM network to the VM.



A DHCP that assigns IPs with an MTU less than 1500 bytes can prevent you configuring the management node network or using the management node UI.

3. Configure the management node network (eth0).



If you need an additional NIC to isolate storage traffic, see instructions on configuring another NIC: [Configure a storage Network Interface Controller \(NIC\)](#).

Configure time sync

1. Ensure time is synced between the management node and the storage cluster using NTP:



Starting with Element 12.3.1, substeps (a) to (e) are performed automatically. For management node 12.3.1, proceed to [substep \(f\)](#) to complete the time sync configuration.

- a. Log in to the management node using SSH or the console provided by your hypervisor.

- b. Stop NTPD:

```
sudo service ntpd stop
```

- c. Edit the NTP configuration file /etc/ntp.conf:

- i. Comment out the default servers (`server 0.gentoo.pool.ntp.org`) by adding a # in front of each.

- ii. Add a new line for each default time server you want to add. The default time servers must be the same NTP servers used on the storage cluster that you will use in a [later step](#).

```
vi /etc/ntp.conf
```

```
#server 0.gentoo.pool.ntp.org
#server 1.gentoo.pool.ntp.org
#server 2.gentoo.pool.ntp.org
#server 3.gentoo.pool.ntp.org
server <insert the hostname or IP address of the default time server>
```

- iii. Save the configuration file when complete.

- d. Force an NTP sync with the newly added server.

```
sudo ntpd -gq
```

- e. Restart NTPD.

```
sudo service ntpd start
```

- f. Disable time synchronization with host via the hypervisor (the following is a VMware example):



If you deploy the mNode in a hypervisor environment other than VMware, for example, from the .iso image in an Openstack environment, refer to the hypervisor documentation for the equivalent commands.

- i. Disable periodic time synchronization:

```
vmware-toolbox-cmd timesync disable
```

- ii. Display and confirm the current status of the service:

```
vmware-toolbox-cmd timesync status
```

- iii. In vSphere, verify that the Synchronize guest time with host box is un-checked in the VM options.



Do not enable this option if you make future changes to the VM.



Do not edit the NTP after you complete the time sync configuration because it affects the NTP when you run the [setup command](#) on the management node.

Set up the management node

1. Configure and run the management node setup command:



You will be prompted to enter passwords in a secure prompt. If your cluster is behind a proxy server, you must configure the proxy settings so you can reach a public network.

```
sudo /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:



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 management virtual IP address (MVIP) of the storage cluster running Element software. Configure the management node with the same storage cluster that you used during [NTP servers configuration](#).
- **--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 Active IQ endpoint parameters to the command:

- **--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. (Recommended): Add the following persistent volume parameters. 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.



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 management virtual IP address (MVIP) 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.

e. (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.

f. Run the `setup-mnode` command.

Configure controller assets

1. Locate the installation ID:

- From a browser, log into the management node REST API UI:
- Go to the storage MVIP and log in. This action causes the certificate to be accepted for the next step.
- Open the inventory service REST API UI on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```

d. Select **Authorize** and complete the following:

- Enter the cluster user name and password.
- Enter the client ID as `mnode-client`.

- iii. Select **Authorize** to begin a session.
 - e. From the REST API UI, select **GET /installations**.
 - f. Select **Try it out**.
 - g. Select **Execute**.
 - h. From the code 200 response body, copy and save the `id` for the installation for use in a later step.
- Your installation has a base asset configuration that was created during installation or upgrade.
2. (NetApp HCI only) Locate the hardware tag for your compute node in vSphere:
 - a. Select the host in the vSphere Web Client navigator.
 - b. Select the **Monitor** tab, and select **Hardware Health**.
 - c. The node BIOS manufacturer and model number are listed. Copy and save the value for `tag` for use in a later step.
 3. Add a vCenter controller asset for NetApp HCI monitoring (NetApp HCI installations only) and Hybrid Cloud Control (for all installations) to the management node known assets:
 - a. Access the mnode service API UI on the management node by entering the management node IP address followed by `/mnode`:

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

- b. Select **Authorize** or any lock icon and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client`.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the window.
- c. Select **POST /assets/{asset_id}/controllers** to add a controller sub-asset.



It is recommended that you create a new NetApp HCC role in vCenter to add a controller sub-asset. This new NetApp HCC role will limit the management node services view to NetApp-only assets. See [Create a NetApp HCC role in vCenter](#).

- d. Select **Try it out**.
- e. Enter the parent base asset ID you copied to your clipboard in the `asset_id` field.
- f. Enter the required payload values with type vCenter and vCenter credentials.
- g. Select **Execute**.

(NetApp HCI only) Configure compute node assets

1. (For NetApp HCI only) Add a compute node asset to the management node known assets:
 - a. Select **POST /assets/{asset_id}/compute-nodes** to add a compute node sub-asset with credentials for the compute node asset.
 - b. Select **Try it out**.

- c. Enter the parent base asset ID you copied to your clipboard in the **asset_id** field.
- d. In the payload, enter the required payload values as defined in the Model tab. Enter **ESXi Host** as type and enter the hardware tag you saved during a previous step for **hardware_tag**.
- e. Select **Execute**.

Find more Information

- [Persistent volumes](#)
- [Add an asset to the management node](#)
- [Configure a storage NIC](#)
- [NetApp Element Plug-in for vCenter Server](#)

Configure a storage Network Interface Controller (NIC)

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

What you'll need

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

Configuration options

Choose the option that is relevant for your environment:

- [Configure a storage Network Interface Controller \(NIC\) for an untagged network interface](#)
- [Configure a storage Network Interface Controller \(NIC\) for a tagged network interface](#)

Configure a storage Network Interface Controller (NIC) for an untagged network interface

Steps

1. Open an SSH or vCenter console.
2. Replace the values in the following command template and run the command:



Values are represented by \$ for each of the required parameters for your new storage network interface. The `cluster` object in the following template is required and can be used for management node host name renaming. `--insecure` or `-k` options should not be used in production environments.

```

curl -u $mnode_user_name:$mnode_password --insecure -X POST \
https://$mnode_IP:442/json-rpc/10.0 \
-H 'Content-Type: application/json' \
-H 'cache-control: no-cache' \
-d '{
  "params": {
    "network": {
      "$eth1": {
        "#default" : false,
        "address" : "$storage_IP",
        "auto" : true,
        "family" : "inet",
        "method" : "static",
        "mtu" : "9000",
        "netmask" : "$subnet_mask",
        "status" : "Up"
      }
    },
    "cluster": {
      "name": "$mnode_host_name"
    }
  },
  "method": "SetConfig"
}
'

```

Configure a storage Network Interface Controller (NIC) for a tagged network interface

Steps

1. Open an SSH or vCenter console.
2. Replace the values in the following command template and run the command:



Values are represented by \$ for each of the required parameters for your new storage network interface. The `cluster` object in the following template is required and can be used for management node host name renaming. `--insecure` or `-k` options should not be used in production environments.

```

curl -u $mnode_user_name:$mnode_password --insecure -X POST \
https://$mnode_IP:442/json-rpc/10.0 \
-H 'Content-Type: application/json' \
-H 'cache-control: no-cache' \
-d ' {
  "params": {
    "network": {
      "$eth1": {
        "#default" : false,
        "address" : "$storage_IP",
        "auto" : true,
        "family" : "inet",
        "method" : "static",
        "mtu" : "9000",
        "netmask" : "$subnet_mask",
        "status" : "Up",
        "virtualNetworkTag" : "$vlan_id"
      }
    },
    "cluster": {
      "name": "$mnode_host_name",
      "cipi": "$eth1.$vlan_id",
      "sipi": "$eth1.$vlan_id"
    }
  },
  "method": "SetConfig"
}
'

```

Find more Information

- [Add an asset to the management node](#)
- [NetApp Element Plug-in for vCenter Server](#)

Recover a management node

You can manually recover and redeploy the management node for your cluster running NetApp Element software if your previous management node used persistent volumes.

You can deploy a new OVA and run a redeploy script to pull configuration data from a previously installed management node running version 11.3 and later.

What you'll need

- Your previous management node was running NetApp Element software version 11.3 or later with [persistent volumes](#) functionality engaged.

- You know the MVIP and SVIP of the cluster containing the persistent volumes.
- Your cluster version is running NetApp Element software 11.3 or later.
- Your installation uses IPv4. The management node 11.3 does not support IPv6.
- You have permission to download software from the NetApp Support Site.
- You have identified the management node image type that is correct for your platform:

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

Steps

1. [Download ISO or OVA and deploy the VM](#)
2. [Configure the network](#)
3. [Configure time sync](#)
4. [Configure the management node](#)

Download ISO or OVA and deploy the VM

1. Download the OVA or ISO for your installation from the [NetApp HCI](#) page on the NetApp Support Site:
 - a. Select **Download Latest Release** and accept the EULA.
 - b. 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 (for example, 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
 - 24GB RAM
 - 400GB virtual disk, thin provisioned
 - One virtual network interface with internet access and access to the storage MVIP.
 - 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.



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.



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.

Configure the network

1. Using the terminal user interface (TUI), create a management node admin user.



To move through the menu options, press the Up or Down arrow keys. To move through the buttons, press Tab. To move from the buttons to the fields, press Tab. To navigate between fields, press the Up or Down arrow keys.

2. Configure the management node network (eth0).



If you need an additional NIC to isolate storage traffic, see instructions on configuring another NIC: [Configure a storage Network Interface Controller \(NIC\)](#).

Configure time sync

1. Ensure time is synced between the management node and the storage cluster using NTP:



Starting with Element 12.3.1, substeps (a) to (e) are performed automatically. For management node 12.3.1, proceed to [substep \(f\)](#) to complete the time sync configuration.

- a. Log in to the management node using SSH or the console provided by your hypervisor.

- b. Stop NTPD:

```
sudo service ntpd stop
```

- c. Edit the NTP configuration file /etc/ntp.conf:

- i. Comment out the default servers (server 0.gentoo.pool.ntp.org) by adding a # in front of each.
- ii. Add a new line for each default time server you want to add. The default time servers must be the same NTP servers used on the storage cluster that you will use in a [later step](#).

```
vi /etc/ntp.conf

#server 0.gentoo.pool.ntp.org
#server 1.gentoo.pool.ntp.org
#server 2.gentoo.pool.ntp.org
#server 3.gentoo.pool.ntp.org
server <insert the hostname or IP address of the default time server>
```

iii. Save the configuration file when complete.

d. Force an NTP sync with the newly added server.

```
sudo ntpd -gg
```

e. Restart NTPD.

```
sudo service ntpd start
```

f. Disable time synchronization with host via the hypervisor (the following is a VMware example):



If you deploy the mNode in a hypervisor environment other than VMware, for example, from the .iso image in an Openstack environment, refer to the hypervisor documentation for the equivalent commands.

i. Disable periodic time synchronization:

```
vmware-toolbox-cmd timesync disable
```

ii. Display and confirm the current status of the service:

```
vmware-toolbox-cmd timesync status
```

iii. In vSphere, verify that the Synchronize guest time with host box is un-checked in the VM options.



Do not enable this option if you make future changes to the VM.



Do not edit the NTP after you complete the time sync configuration because it affects the NTP when you run the [redeploy command](#) on the management node.

Configure the management node

1. Create a temporary destination directory for the management services bundle contents:

```
mkdir -p /sf/etc/mnode/mnode-archive
```

2. Download the management services bundle (version 2.15.28 or later) that was previously installed on the existing management node and save it in the /sf/etc/mnode/ directory.
3. Extract the downloaded bundle using the following command, replacing the value in [] brackets (including the brackets) with the name of the bundle file:

```
tar -C /sf/etc/mnode -xvf /sf/etc/mnode/[management services bundle file]
```

4. Extract the resulting file to the /sf/etc/mnode-archive directory:

```
tar -C /sf/etc/mnode/mnode-archive -xvf /sf/etc/mnode/services_deploy_bundle.tar.gz
```

5. Create a configuration file for accounts and volumes:

```
echo '{"trident": true, "mvip": "[mvip IP address]", "account_name": "[persistent volume account name]"}' | sudo tee /sf/etc/mnode/mnode-archive/management-services-metadata.json
```

- a. Replace the value in [] brackets (including the brackets) for each of the following required parameters:
 - **[mvip IP address]**: The management virtual IP address of the storage cluster. Configure the management node with the same storage cluster that you used during [NTP servers configuration](#).
 - **[persistent volume account name]**: The name of the account associated with all persistent volumes in this storage cluster.
6. Configure and run the management node redeploy command to connect to persistent volumes hosted on the cluster and start services with previous management node configuration data:



You will be prompted to enter passwords in a secure prompt. If your cluster is behind a proxy server, you must configure the proxy settings so you can reach a public network.

```
sudo /sf/packages/mnode/redeploy-mnode --mnode_admin_user [username]
```

- a. Replace the value in [] brackets (including the brackets) with the user name 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.



You can add the user name or allow the script to prompt you for the information.

- b. Run the redeploy-mnode command. The script displays a success message when the redeployment is complete.

- c. If you access Element or NetApp HCI web interfaces (such as the management node or NetApp Hybrid Cloud Control) using the Fully Qualified Domain Name (FQDN) of the system, [reconfigure authentication for the management node](#).



If you had previously disabled SSH functionality on the management node, you need to [disable SSH again](#) on the recovered management node. SSH capability that provides [NetApp Support remote support tunnel \(RST\) session access](#) is enabled on the management node by default.

Find more Information

- [Persistent volumes](#)
- [NetApp Element Plug-in for vCenter Server](#)

Access the management node

Beginning with NetApp Element software version 11.3, the management node contains two UIs: a UI for managing REST-based services and a per-node UI for managing network and cluster settings and operating system tests and utilities.

For clusters running Element software version 11.3 or later, you can make use one of two interfaces:

- By using the management node UI (`https:// [mNode IP]:442`), you can make changes to network and cluster settings, run system tests, or use system utilities.
- By using the built-in REST API UI (`https:// [mNode IP] /mnode`), you can run or understand APIs relating to the management node services, including proxy server configuration, service level updates, or asset management.

Access the management node per-node UI

From the per-node UI, you can access network and cluster settings and utilize system tests and utilities.

Steps

1. Access the per-node UI for the management node by entering the management node IP address followed by :442

```
https:// [IP address]:442
```

Management

Network Settings - Management

Method : static

Link Speed : 1000

IPv4 Address : 10.117.148.201

IPv4 Subnet Mask : 255.255.192.0

IPv4 Gateway Address : 10.117.192.254

IPv6 Address :

IPv6 Gateway Address :

MTU : 1500

DNS Servers : 10.117.20.40, 10.116.133.40

Search Domains : dan solidfire.net, one.dan.solidfire

Status : UpAndRunning ▾

Routes

Add

Reset Changes | Save Changes

2. Enter the management node user name and password when prompted.

Access the management node REST API UI

From the REST API UI, you can access a menu of service-related APIs that control management services on the management node.

Steps

1. To access the REST API UI for management services, enter the management node IP address followed by /mnode:

```
https://[IP address]/mnode
```

MANAGEMENT SERVICES API^{1.0}

[Base URL: /node]
https://10.117.1.100/mnode/swagger.json

The configuration REST service for MANAGEMENT SERVICES
NetApp - Website
NetApp Commercial Software License

Authorize 

logs Log service

assets Asset service

Method	Path	Description	Lock
GET	/logs	Get logs from the MNODE service(s)	
POST	/assets	Add a new asset	
GET	/assets	Get all assets	
GET	/assets/compute-nodes	Get all compute nodes	
GET	/assets/compute-nodes/{compute_node_id}	Get a specific compute node by ID	
GET	/assets/controllers	Get all controllers	
GET	/assets/controllers/{controller_id}	Get a specific controller by ID	
GET	/assets/storage-clusters	Get all storage clusters	
GET	/assets/storage-clusters/{storage_cluster_id}	Get a specific storage cluster by ID	
PUT	/assets/{asset_id}	Modify an asset with a specific ID	
DELETE	/assets/{asset_id}	Delete an asset with a specific ID	
GET	/assets/{asset_id}	Get an asset by it's ID	
POST	/assets/{asset_id}/compute-nodes	Add a compute asset	
GET	/assets/{asset_id}/compute-nodes	Get compute assets	
PUT	/assets/{asset_id}/compute-nodes/{compute_id}	Update a specific compute node asset	
DELETE	/assets/{asset_id}/compute-nodes/{compute_id}	Delete a specific compute node asset	

2. Select **Authorize** or any lock icon and enter cluster admin credentials for permissions to use APIs.

Find more Information

- Enable Active IQ and NetApp HCI monitoring
- NetApp Element Plug-in for vCenter Server

Change the management node default SSL certificate

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

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

You can use the following API methods to get more information about the default SSL certificate and make

changes.

- **GetNodeSSLCertificate**

You can use the [GetNodeSSLCertificate method](#) to retrieve information about the currently installed SSL certificate including all certificate details.

- **SetNodeSSLCertificate**

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

- **RemoveNodeSSLCertificate**

This [RemoveNodeSSLCertificate method](#) removes the currently installed SSL certificate and private key. The cluster then generates a new self-signed certificate and private key.

Find more information

- [Change the Element software default SSL certificate](#)
- [What are the requirements around setting custom SSL certificates in Element Software?](#)
- [SolidFire and Element Software Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

Work with the management node UI

Management node UI overview

With the management node UI (<https://<ManagementNodeIP>:442>), you can make changes to network and cluster settings, run system tests, or use system utilities.

Tasks you can perform with the management node UI:

- [Configure alert monitoring on NetApp HCI](#)
- [Modify and test the management node network, cluster, and system settings](#)
- [Run system utilities from the management node](#)

Find more information

- [Access the management node](#)
- [NetApp Element Plug-in for vCenter Server](#)

Configure alert monitoring on NetApp HCI

You can configure settings to monitor alerts on your NetApp HCI system.

NetApp HCI alert monitoring forwards NetApp HCI storage cluster system alerts to vCenter Server, enabling you to view all alerts for NetApp HCI from the vSphere Web Client interface.

1. Open the per-node management node UI ([https://\[IP address\]:442](https://[IP address]:442)).

2. Select the **Alert Monitor** tab.
3. Configure the alert monitoring options.

Alert monitoring options

options	Description
Run Alert Monitor Tests	<p>Runs the monitor system tests to check for the following:</p> <ul style="list-style-type: none"> • NetApp HCI and VMware vCenter connectivity • Pairing of NetApp HCI and VMware vCenter through datastore information supplied by the QoSIOC service • Current NetApp HCI alarm and vCenter alarm lists
Collect Alerts	<p>Enables or disables the forwarding of NetApp HCI storage alarms to vCenter. You can select the target storage cluster from the drop-down list. The default setting for this option is Enabled.</p>
Collect Best Practice Alerts	<p>Enables or disables the forwarding of NetApp HCI storage Best Practice alerts to vCenter. Best Practice alerts are faults that are triggered by a sub-optimal system configuration. The default setting for this option is Disabled. When disabled, NetApp HCI storage Best Practice alerts do not appear in vCenter.</p>
Send Support Data To AIQ	<p>Controls the flow of support and monitoring data from VMware vCenter to NetApp SolidFire Active IQ.</p> <p>Options are the following:</p> <ul style="list-style-type: none"> • Enabled: All vCenter alarms, NetApp HCI storage alarms, and support data are sent to NetApp SolidFire Active IQ. This enables NetApp to proactively support and monitor the NetApp HCI installation, so that possible problems can be detected and resolved before affecting the system. • Disabled: No vCenter alarms, NetApp HCI storage alarms, or support data are sent to NetApp SolidFire Active IQ. <p> If you turned off the Send data to AIQ option using NetApp Deployment Engine, you need to enable telemetry again using the management node REST API to configure the service from this page.</p>

options	Description
Send Compute Node Data To AIQ	<p>Controls the flow of support and monitoring data from the compute nodes to NetApp SolidFire Active IQ.</p> <p>Options are the following:</p> <ul style="list-style-type: none"> Enabled: Support and monitoring data about the compute nodes is transmitted to NetApp SolidFire Active IQ to enable proactive support for the compute node hardware. Disabled: Support and monitoring data about the compute nodes is not transmitted to NetApp SolidFire Active IQ. <p> If you turned off the Send data to AIQ option using NetApp Deployment Engine, you need to enable telemetry again using the management node REST API to configure the service from this page.</p>

Find more Information

- [NetApp Element Plug-in for vCenter Server](#)

Modify and test the management node network, cluster, and system settings

You can modify and test the management node network, cluster, and system settings.

- [Update management node network settings](#)
- [Update management node cluster settings](#)
- [Test the management node settings](#)

Update management node network settings

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

1. Open the per-node management node UI.
2. Select the **Network Settings** tab.
3. View or enter the following information:
 - a. **Method:** Choose one of the following methods to configure the interface:
 - **loopback:** Use to define the IPv4 loopback interface.
 - **manual:** Use to define interfaces for which no configuration is done by default.
 - **dhop:** Use to obtain an IP address via DHCP.
 - **static:** Use to define Ethernet interfaces with statically allocated IPv4 addresses.

- b. **Link Speed**: The speed negotiated by the virtual NIC.
- c. **IPv4 Address**: The IPv4 address for the eth0 network.
- d. **IPv4 Subnet Mask**: Address subdivisions of the IPv4 network.
- e. **IPv4 Gateway Address**: Router network address to send packets out of the local network.
- f. **IPv6 Address**: The IPv6 address for the eth0 network.
- g. **IPv6 Gateway Address**: Router network address to send packets out of the local network.



The IPv6 options are not supported for 11.3 or later versions of the management node.

- h. **MTU**: Largest packet size that a network protocol can transmit. Must be greater than or equal to 1500. If you add a second storage NIC, the value should be 9000.
- i. **DNS Servers**: Network interface used for cluster communication.
- j. **Search Domains**: Search for additional MAC addresses available to the system.
- k. **Status**: Possible values:
 - UpAndRunning
 - Down
 - Up
- l. **Routes**: Static routes to specific hosts or networks via the associated interface the routes are configured to use.

Update management node cluster settings

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

1. Open the per-node management node UI.
2. Select the **Cluster Settings** tab.
3. View or enter the following information:
 - **Role**: Role the management node has in the cluster. Possible value: Management.
 - **Version**: Element software version running on the cluster.
 - **Default Interface**: Default network interface used for management node communication with the cluster running Element software.

Test the management node settings

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

1. Open the per-node management node UI.
2. In the management node UI, select **System Tests**.
3. Complete any of the following:
 - a. To verify that the network settings you configured are valid for the system, select **Test Network Config**.
 - b. To test network connectivity to all nodes in the cluster on both 1G and 10G interfaces using ICMP packets, select **Test Ping**.

4. View or enter the following:

- **Hosts:** Specify a comma-separated list of addresses or host names of devices to ping.
- **Attempts:** Specify the number of times the system should repeat the test ping. Default: 5.
- **Packet Size:** Specify the number of bytes to send in the ICMP packet that is sent to each IP. The number of bytes must be less than the maximum MTU specified in the network configuration.
- **Timeout mSec:** Specify the number of milliseconds to wait for each individual ping response. Default: 500 ms.
- **Total Timeout Sec:** Specify the time in seconds the ping should wait for a system response before issuing the next ping attempt or ending the process. Default: 5.
- **Prohibit Fragmentation:** Enable the DF (do not fragment) flag for the ICMP packets.

Find more Information

- [NetApp Element Plug-in for vCenter Server](#)

Run system utilities from the management node

You can use the per-node UI for the management node to create or delete cluster support bundles, reset node configuration settings, or restart networking.

Steps

1. Open the per-node management node UI using the management node admin credentials.
2. Select **System Utilities**.
3. Select the button for the utility that you want to run:
 - a. **Control Power:** Reboots, power cycles, or shuts down the node. Specify any of the following options.



This operation causes temporary loss of networking connectivity.

- **Action:** Options include Restart and Halt (power off).
 - **Wakeup Delay:** Any additional time before the node comes back online.
- b. **Create Cluster Support Bundle:** Creates the cluster support bundle to assist NetApp Support diagnostic evaluations of one or more nodes in a cluster. Specify the following options:
 - **Bundle Name:** Unique name for each support bundle created. If no name is provided, then "supportbundle" and the node name are used as the file name.
 - **Mvip:** The MVIP of the cluster. Bundles are gathered from all nodes in the cluster. This parameter is required if the Nodes parameter is not specified.
 - **Nodes:** The IP addresses of the nodes from which to gather bundles. Use either Nodes or Mvip, but not both, to specify the nodes from which to gather bundles. This parameter is required if Mvip is not specified.
 - **Username:** The cluster admin user name.
 - **Password:** The cluster admin password.
 - **Allow Incomplete:** Allows the script to continue to run if bundles cannot be gathered from one or more of the nodes.
 - **Extra Args:** This parameter is fed to the `sf_make_support_bundle` script. This parameter should be used only at the request of NetApp Support.

- c. **Delete All Support Bundles:** Deletes any current support bundles on the management node.
- d. **Reset Node:** Resets the management node to a new install image. This changes all settings except the network configuration to the default state. Specify the following options:
 - **Build:** The URL to a remote Element software image to which the node will be reset.
 - **Options:** Specifications for running the reset operations. Details are provided by NetApp Support, if required.



This operation causes temporary loss of networking connectivity.

- e. **Restart Networking:** Restarts all networking services on the management node.



This operation causes temporary loss of networking connectivity.

Find more Information

- [NetApp Element Plug-in for vCenter Server](#)

Work with the management node REST API

Management node REST API UI overview

By using the built-in REST API UI (<https://<ManagementNodeIP>/mnode>), you can run or understand APIs relating to the management node services, including proxy server configuration, service level updates, or asset management.

Tasks you can perform with REST APIs:

Authorization

- [Get authorization to use REST APIs](#)

Asset configuration

- [Enable Active IQ and NetApp HCI monitoring](#)
- [Configure a proxy server for the management node](#)
- [Configure NetApp Hybrid Cloud Control for multiple vCenters](#)
- [Add compute and controller assets to the management node](#)
- [Create and manage storage cluster assets](#)

Asset management

- [View or edit existing controller assets](#)
- [Create and manage storage cluster assets](#)
- [Remove an asset from the management node](#)
- [Use the REST API to collect NetApp HCI logs](#)
- [Verify management node OS and services versions](#)

- [Getting logs from management services](#)

Find more information

- [Access the management node](#)
- [NetApp Element Plug-in for vCenter Server](#)

Get authorization to use REST APIs

You must authorize before you can use APIs for management services in the REST API UI. You do this by obtaining an access token.

To obtain a token, you provide cluster admin credentials and a client ID. Each token lasts approximately ten minutes. After a token expires, you can authorize again for a new access token.

Authorization functionality is set up for you during management node installation and deployment. The token service is based on the storage cluster you defined during setup.

What you'll need

- Your cluster version should be running NetApp Element software 11.3 or later.
- You should have deployed a management node running version 11.3 or later.

API command

```
TOKEN=`curl -k -X POST https://MVIP/auth/connect/token -F client_id=mnode-client -F grant_type=password -F username=CLUSTER_ADMIN -F password=CLUSTER_PASSWORD|awk -F':|'{print $2}'|awk -F',' '{print $1}'|sed s/\"//g`
```

REST API UI steps

1. Access the REST API UI for the service by entering the management node IP address followed by the service name, for example /mnode/:

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

2. Select **Authorize**.



Alternately, you can select the lock icon next to any service API.

3. Complete the following:
 - a. Enter the cluster user name and password.
 - b. Enter the client ID as mnode-client.
 - c. Do not enter a value for the client secret.
 - d. Select **Authorize** to begin a session.
4. Close the **Available authorizations** dialog box.



If you try to run a command after the token expires, a 401 Error: UNAUTHORIZED message appears. If you see this, authorize again.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Enable Active IQ and NetApp HCI monitoring

You can enable Active IQ storage monitoring for NetApp HCI and NetApp HCI compute monitoring if you did not already do so during installation or upgrade. You might need to use this procedure if you disabled telemetry using the NetApp HCI Deployment Engine.

The Active IQ collector service forwards configuration data and Element software-based cluster performance metrics to NetApp Active IQ for historical reporting and near real-time performance monitoring. The NetApp HCI monitoring service enables forwarding of storage cluster faults to vCenter for alert notification.

What you'll need

- Your storage cluster is running NetApp Element software 11.3 or later.
- You have deployed a management node running version 11.3 or later.
- You have internet access. The Active IQ collector service cannot be used from dark sites.

Steps

1. Get the base asset ID for the installation:
 - a. Open the inventory service REST API UI on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```

- b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client`.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the window.
 - c. From the REST API UI, select **GET /installations**.
 - d. Select **Try it out**.
 - e. Select **Execute**.
 - f. From the code 200 response body, copy the `id` for the installation.

```
{
  "installations": [
    {
      "_links": {
        "collection": "https://10.111.211.111/inventory/1/installations",
        "self": "https://10.111.217.111/inventory/1/installations/abcd01e2-ab00-1xxx-91ee-12f111xxc7x0x"
      },
      "id": "abcd01e2-ab00-1xxx-91ee-12f111xxc7x0x",
    }
  ]
}
```



Your installation has a base asset configuration that was created during installation or upgrade.

2. Activate telemetry:

- Access the mnode service API UI on the management node by entering the management node IP address followed by /mnode:

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

- Select **Authorize** or any lock icon and complete the following:

- Enter the cluster user name and password.
- Enter the client ID as mnode-client.
- Select **Authorize** to begin a session.
- Close the window.

- Configure the base asset:

- Select **PUT /assets/{asset_id}**.
- Select **Try it out**.
- Enter the following in the JSON payload:

```
{
  "telemetry_active": true
  "config": {}
}
```

- Enter the base ID from the previous step in **asset_ID**.
- Select **Execute**.

The Active IQ service is automatically restarted whenever assets are changed. Modifying assets results in a short delay before settings are applied.

3. If you have not already done so, add a vCenter controller asset for NetApp HCI monitoring (NetApp HCI installations only) and Hybrid Cloud Control (for all installations) to the management node known assets:



A controller asset is required for NetApp HCI monitoring services.

- a. Select **POST /assets/{asset_id}/controllers** to add a controller sub-asset.
- b. Select **Try it out**.
- c. Enter the parent base asset ID you copied to your clipboard in the **asset_id** field.
- d. Enter the required payload values with **type** as vCenter and vCenter credentials.

```
{  
  "username": "string",  
  "password": "string",  
  "ip": "string",  
  "type": "vCenter",  
  "host_name": "string",  
  "config": {}  
}
```



ip is the vCenter IP address.

- e. Select **Execute**.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Configure NetApp Hybrid Cloud Control for multiple vCenters

You can configure NetApp Hybrid Cloud Control to manage assets from two or more vCenters that are not using Linked Mode.

You should use this process after your initial installation when you need to add assets for a recently scaled installation or when new assets were not added automatically to your configuration. Use these APIs to add assets that are recent additions to your installation.

What you'll need

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

Steps

1. [Add new vCenters as controller assets](#) to the management node configuration.
2. [Add new compute nodes as compute assets](#) to the management node configuration.



You might need to [change BMC credentials for compute nodes](#) to resolve a `Hardware ID not available` or `Unable to Detect` error indicated in NetApp Hybrid Cloud Control.

3. Refresh the inventory service API on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```



As an alternative, you can wait 15 minutes for the inventory to update in NetApp Hybrid Cloud Control UI.

a. Select **Authorize** and complete the following:

- i. Enter the cluster user name and password.
- ii. Enter the client ID as `mnode-client`.
- iii. Select **Authorize** to begin a session.
- iv. Close the window.

b. From the REST API UI, select **GET /installations**.

c. Select **Try it out**.

d. Select **Execute**.

e. From the response, copy the installation asset ID ("`id`").

f. From the REST API UI, select **GET /installations/{id}**.

g. Select **Try it out**.

h. Set refresh to `True`.

i. Paste the installation asset ID into the **id** field.

j. Select **Execute**.

4. Refresh the NetApp Hybrid Cloud Control browser to see the changes.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Add compute and controller assets to the management node

You can add compute and controller assets to the management node configuration using the REST API UI.

You might need to add an asset if you recently scaled your installation and new assets were not added automatically to your configuration. Use these APIs to add assets that are recent additions to your installation.

What you'll need

- Your cluster version is running NetApp Element software 11.3 or later.
- You have deployed a management node running version 11.3 or later.
- You have [created a new NetApp HCC role in vCenter](#) to limit the management node services view to NetApp-only assets.
- You have the vCenter management IP address and credentials.
- You have the compute node (ESXi) management IP address and root credentials.

- You have the hardware (BMC) management IP address and administrator credentials.

About this task

(NetApp HCI only) If you do not see compute nodes in Hybrid Cloud Control (HCC) after scaling your NetApp HCI system, you can add a compute node using the `POST /assets/{asset_id}/compute-nodes` described in this procedure.

Steps

1. Get the base asset ID for the installation:
 - a. Open the inventory service REST API UI on the management node:


```
https://<ManagementNodeIP>/inventory/1/
```
 - b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client`.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the window.
 - c. From the REST API UI, select **GET /installations**.
 - d. Select **Try it out**.
 - e. Select **Execute**.
 - f. From the code 200 response body, copy the `id` for the installation.

```
{
  "installations": [
    {
      "_links": {
        "collection": "https://10.111.211.111/inventory/1/installations",
        "self": "https://10.111.217.111/inventory/1/installations/abcd01e2-ab00-1xxx-91ee-12f111xxc7x0x"
      },
      "id": "abcd01e2-ab00-1xxx-91ee-12f111xxc7x0x",
    }
  ]
}
```



Your installation has a base asset configuration that was created during installation or upgrade.

- g. From the REST API UI, select **GET /installations/{id}**.
- h. Select **Try it out**.
- i. Paste the installation asset ID into the `id` field.
- j. Select **Execute**.

- k. From the response, copy and save the cluster controller ID ("controllerId") for use in a later step.
2. (For compute nodes only) [Locate the hardware tag for your compute node](#) in vSphere.
3. To add a controller asset (vCenter), compute node (ESXi), or hardware (BMC) to an existing base asset, select one of the following.

Option	Description
POST /assets/{asset_id}/controllers	<p>1. Open the mNode service REST API UI on the management node:</p> <div style="border: 1px solid #ccc; padding: 5px; margin-top: 10px; text-align: center;"> <code>https://<ManagementNodeIP>/mnode</code> </div> <p>a. Select Authorize and complete the following:</p> <ol style="list-style-type: none"> Enter the cluster user name and password. Enter the client ID as <code>mnode-client</code>. Select Authorize to begin a session. Close the window. <p>2. Select POST /assets/{asset_id}/controllers.</p> <p>3. Select Try it out.</p> <p>4. Enter the parent base asset ID in the asset_id field.</p> <p>5. Add the required values to the payload.</p> <p>6. Select Execute.</p>

Option	Description
POST /assets/{asset_id}/compute-nodes	<p>1. Open the mNode service REST API UI on the management node:</p> <div data-bbox="763 255 1494 361" style="border: 1px solid #ccc; padding: 5px; border-radius: 5px; margin: 10px 0;"> <p style="margin: 0;">https://<ManagementNodeIP>/mnode</p> </div> <p>a. Select Authorize and complete the following:</p> <ol style="list-style-type: none"> Enter the cluster user name and password. Enter the client ID as mnode-client. Select Authorize to begin a session. Close the window. <p>2. Select POST /assets/{asset_id}/compute-nodes.</p> <p>3. Select Try it out.</p> <p>4. Enter the parent base asset ID you copied in an earlier step in the asset_id field.</p> <p>5. In the payload, do the following:</p> <ol style="list-style-type: none"> Enter the management IP for the node in the ip field. For hardwareTag, enter the hardware tag value you saved in an earlier step. Enter other values, as required. <p>6. Select Execute.</p>
POST /assets/{asset_id}/hardware-nodes	<p>1. Open the mNode service REST API UI on the management node:</p> <div data-bbox="763 1262 1494 1368" style="border: 1px solid #ccc; padding: 5px; border-radius: 5px; margin: 10px 0;"> <p style="margin: 0;">https://<ManagementNodeIP>/mnode</p> </div> <p>a. Select Authorize and complete the following:</p> <ol style="list-style-type: none"> Enter the cluster user name and password. Enter the client ID as mnode-client. Select Authorize to begin a session. Close the window. <p>2. Select POST /assets/{asset_id}/hardware-nodes.</p> <p>3. Select Try it out.</p> <p>4. Enter the parent base asset ID in the asset_id field.</p> <p>5. Add the required values to the payload.</p> <p>6. Select Execute.</p>

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

How to locate a hardware tag for a compute node

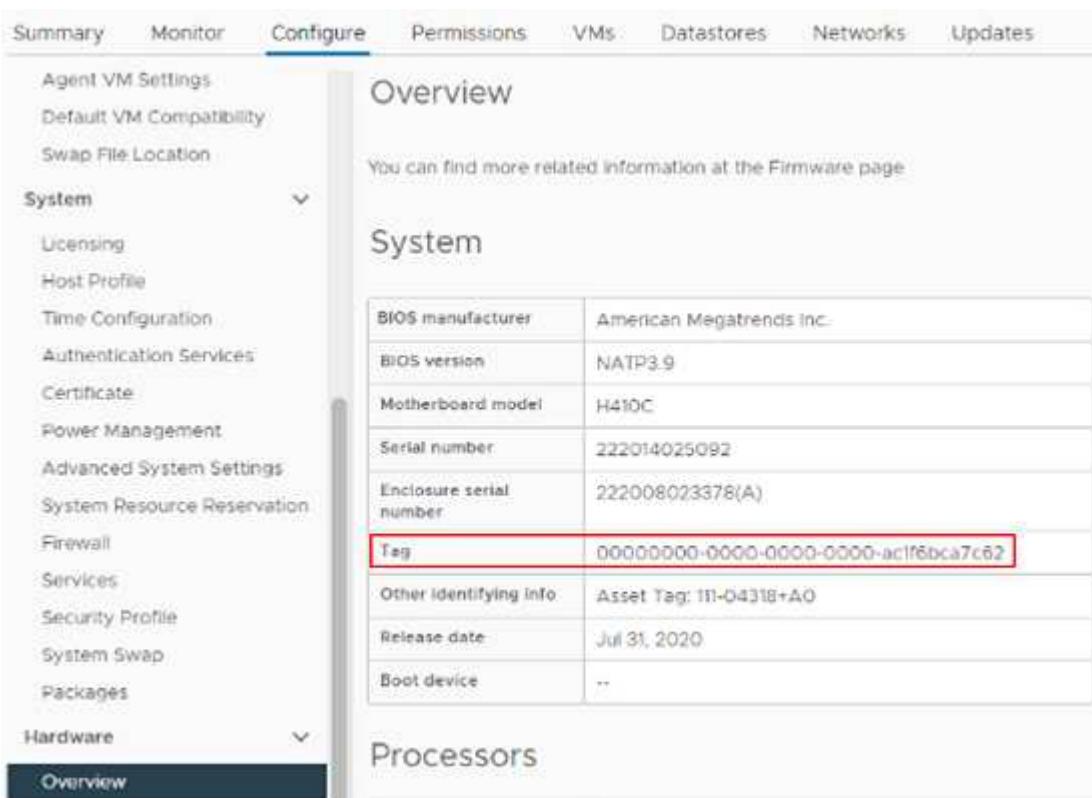
You require the hardware tag to add your compute node assets to the management node configuration using the REST API UI.

VMware vSphere 7.0

Locate the hardware tag for a compute node in VMware vSphere Web Client 7.0.

Steps

1. Select the host in the vSphere Web Client navigator.
2. Select the **Configure** tab.
3. From the sidebar, select **Hardware > Overview**. Check if the hardware tag is listed in the **System** table.



The screenshot shows the 'Configure' tab selected in the top navigation bar. The sidebar on the left is expanded to show the 'System' section, which includes options like Licensing, Host Profile, Time Configuration, Authentication Services, Certificate, Power Management, Advanced System Settings, System Resource Reservation, Firewall, Services, Security Profile, System Swap, and Packages. Below the sidebar, the 'Hardware' section is expanded, and 'Overview' is selected. The main content area is titled 'Overview' and contains a message: 'You can find more related information at the Firmware page'. Below this is a table titled 'System' with the following data:

BIOS manufacturer	American Megatrends Inc.
BIOS version	NATP3.9
Motherboard model	H410C
Serial number	222014025092
Enclosure serial number	222008023378(A)
Tag	00000000-0000-0000-0000-ac1f60ca7c62
Other identifying info	Asset Tag: 111-04318+A0
Release date	Jul 31, 2020
Boot device	--

4. Copy and save the value for **Tag**.

[5. Add your compute and controller assets to the management node.](#)

VMware vSphere 6.7 and 6.5

Locate the hardware tag for a compute node in VMware vSphere Web Client 6.7 and 6.5.

Steps

1. Select the host in the vSphere Web Client navigator.
2. Select the **Monitor** tab, and select **Hardware Health**.
3. Check if the tag is listed with the BIOS manufacturer and model number.

4. Copy and save the value for **Tag**.

5. [Add your compute and controller assets to the management node.](#)

Create and manage storage cluster assets

You can add new storage cluster assets to the management node, edit the stored credentials for known storage cluster assets, and delete storage cluster assets from the management node using the REST API.

What you'll need

- Ensure that your storage cluster version is running NetApp Element software 11.3 or later.
- Ensure that you have deployed a management node running version 11.3 or later.

Storage cluster asset management options

Choose one of the following options:

- [Retrieve the installation ID and cluster ID of a storage cluster asset](#)
- [Add a new storage cluster asset](#)
- [Edit the stored credentials for a storage cluster asset](#)
- [Delete a storage cluster asset](#)

Retrieve the installation ID and cluster ID of a storage cluster asset

You can use the REST API get the installation ID and the ID of the storage cluster. You need the installation ID to add a new storage cluster asset, and the cluster ID to modify or delete a specific storage cluster asset.

Steps

1. Access the REST API UI for the inventory service by entering the management node IP address followed by `/inventory/1/`:

```
https://<ManagementNodeIP>/inventory/1/
```

2. Select **Authorize** or any lock icon and complete the following:

- a. Enter the cluster user name and password.
 - b. Enter the client ID as `mnode-client`.
 - c. Select **Authorize** to begin a session.
 - d. Close the window.
3. Select **GET /installations**.
4. Select **Try it out**.
5. Select **Execute**.

The API returns a list of all known installations.

6. From the code 200 response body, save the value in the `id` field, which you can find in the list of installations. This is the installation ID. For example:

```
"installations": [
  {
    "id": "1234a678-12ab-35dc-7b4a-1234a5b6a7ba",
    "name": "my-hci-installation",
    "_links": {
      "collection": "https://localhost/inventory/1/installations",
      "self": "https://localhost/inventory/1/installations/1234a678-
12ab-35dc-7b4a-1234a5b6a7ba"
    }
  }
]
```

7. Access the REST API UI for the storage service by entering the management node IP address followed by `/storage/1/`:

```
https://<ManagementNodeIP>/storage/1/
```

8. Select **Authorize** or any lock icon and complete the following:
- a. Enter the cluster user name and password.
 - b. Enter the client ID as `mnode-client`.
 - c. Select **Authorize** to begin a session.
 - d. Close the window.
9. Select **GET /clusters**.
10. Select **Try it out**.
11. Enter the installation ID you saved earlier into the `installationId` parameter.
12. Select **Execute**.

The API returns a list of all known storage clusters in this installation.

13. From the code 200 response body, find the correct storage cluster and save the value in the cluster's `storageId` field. This is the storage cluster ID.

Add a new storage cluster asset

You can use the REST API to add one or more new storage cluster assets to the management node inventory. When you add a new storage cluster asset, it is automatically registered with the management node.

What you'll need

- You have copied the [storage cluster ID and installation ID](#) for any storage clusters you want to add.
- If you are adding more than one storage node, you have read and understood the limitations of the [authoritative cluster](#) and multiple storage cluster support.



All users defined on the authoritative cluster are defined as users on all other clusters tied to the Hybrid Cloud Control instance.

Steps

1. Access the REST API UI for the storage service by entering the management node IP address followed by `/storage/1/`:

```
https://<ManagementNodeIP>/storage/1/
```

2. Select **Authorize** or any lock icon and complete the following:

- a. Enter the cluster user name and password.
- b. Enter the client ID as `mnode-client`.
- c. Select **Authorize** to begin a session.
- d. Close the window.

3. Select **POST /clusters**.

4. Select **Try it out**.

5. Enter the new storage cluster's information in the following parameters in the **Request body** field:

```
{
  "installationId": "a1b2c34d-e56f-1a2b-c123-1ab2cd345d6e",
  "mvip": "10.0.0.1",
  "password": "admin",
  "userId": "admin"
}
```

Parameter	Type	Description
<code>installationId</code>	string	The installation in which to add the new storage cluster. Enter the installation ID you saved earlier into this parameter.

Parameter	Type	Description
mvip	string	The IPv4 management virtual IP address (MVIP) of the storage cluster.
password	string	The password used to communicate with the storage cluster.
userId	string	The user ID used to communicate with the storage cluster (the user must have administrator privileges).

6. Select **Execute**.

The API returns an object containing information about the newly added storage cluster asset, such as the name, version, and IP address information.

Edit the stored credentials for a storage cluster asset

You can edit the stored credentials that the management node uses to log in to a storage cluster. The user you choose must have cluster admin access.



Ensure you have followed the steps in [Retrieve the installation ID and cluster ID of a storage cluster asset](#) before continuing.

Steps

1. Access the REST API UI for the storage service by entering the management node IP address followed by `/storage/1/`:

```
https://<ManagementNodeIP>/storage/1/
```

2. Select **Authorize** or any lock icon and complete the following:

- a. Enter the cluster user name and password.
- b. Enter the client ID as `mnode-client`.
- c. Select **Authorize** to begin a session.
- d. Close the window.

3. Select **PUT /clusters/{storageId}**.

4. Select **Try it out**.

5. Paste the storage cluster ID you copied earlier into the `storageId` parameter.

6. Change one or both of the following parameters in the **Request body** field:

```
{
  "password": "adminadmin",
  "userId": "admin"
}
```

Parameter	Type	Description
password	string	The password used to communicate with the storage cluster.
userId	string	The user ID used to communicate with the storage cluster (the user must have administrator privileges).

7. Select **Execute**.

Delete a storage cluster asset

You can delete a storage cluster asset if the storage cluster is no longer in service. When you remove a storage cluster asset, it is automatically unregistered from the management node.



Ensure you have followed the steps in [Retrieve the installation ID and cluster ID of a storage cluster asset](#) before continuing.

Steps

1. Access the REST API UI for the storage service by entering the management node IP address followed by `/storage/1/`:

```
https://<ManagementNodeIP>/storage/1/
```

2. Select **Authorize** or any lock icon and complete the following:

- a. Enter the cluster user name and password.
- b. Enter the client ID as `mnode-client`.
- c. Select **Authorize** to begin a session.
- d. Close the window.

3. Select **DELETE /clusters/{storageId}**.

4. Select **Try it out**.

5. Enter the storage cluster ID you copied earlier in the `storageId` parameter.

6. Select **Execute**.

Upon success, the API returns an empty response.

Find more information

- [Authoritative cluster](#)
- [NetApp Element Plug-in for vCenter Server](#)

View or edit existing controller assets

You can view information about and edit existing VMware vCenter controllers in the management node configuration using the REST API. Controllers are VMware vCenter instances registered to the management node for your NetApp HCI installation.

What you'll need

- Ensure that your cluster version is running NetApp Element software 11.3 or later.
- Ensure that you have deployed a management node running version 11.3 or later.

Access the management services REST API

Steps

1. Access the REST API UI for management services by entering the management node IP address followed by `/vcenter/1/`:

```
https://<ManagementNodeIP>/vcenter/1/
```

2. Select **Authorize** or any lock icon and complete the following:
 - a. Enter the cluster user name and password.
 - b. Enter the client ID as `mnode-client`.
 - c. Select **Authorize** to begin a session.
 - d. Close the window.

View stored information about existing controllers

You can list existing vCenter controllers that are registered with the management node and view stored information about them using the REST API.

Steps

1. Select **GET /compute/controllers**.
2. Select **Try it out**.
3. Select **Execute**.

The API returns a list of all known vCenter controllers, along with the IP address, controller ID, hostname, and user ID used to communicate with each controller.

4. If you want the connection status of a particular controller, copy the controller ID from the `id` field of that controller to your clipboard and see [View the status of an existing controller](#).

View the status of an existing controller

You can view the status of any of the existing vCenter controllers registered with the management node. The API returns a status indicating whether NetApp Hybrid Cloud Control can connect with the vCenter controller as well as the reason for that status.

Steps

1. Select **GET /compute/controllers/{controller_id}/status**.
2. Select **Try it out**.
3. Enter the controller ID you copied earlier in the `controller_id` parameter.
4. Select **Execute**.

The API returns a status of this particular vCenter controller, along with a reason for that status.

Edit the stored properties of a controller

You can edit the stored user name or password for any of the existing vCenter controllers registered with the management node. You cannot edit the stored IP address of an existing vCenter controller.

Steps

1. Select **PUT /compute/controllers/{controller_id}**.
2. Enter the controller ID of a vCenter controller in the `controller_id` parameter.
3. Select **Try it out**.
4. Change either of the following parameters in the **Request body** field:

Parameter	Type	Description
<code>userId</code>	string	Change the user ID used to communicate with the vCenter controller (the user must have administrator privileges).
<code>password</code>	string	Change the password used to communicate with the vCenter controller.

5. Select **Execute**.

The API returns updated controller information.

Find more information

- [Add an asset to the management node](#)
- [NetApp Element Plug-in for vCenter Server](#)

Remove an asset from the management node

If you physically replace a compute node or need to remove it from the NetApp HCI cluster, you must remove the compute node asset using the management node APIs.

What you'll need

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

Steps

1. Enter the management node IP address followed by /mnode/1/:

```
https://<ManagementNodeIP>/mnode/1/
```

2. Select **Authorize** or any lock icon and enter cluster admin credentials for permissions to use APIs.

- a. Enter the cluster user name and password.
- b. Select **Request body** from the type drop-down list if the value is not already selected.
- c. Enter the client ID as mnode-client if the value is not already populated.
- d. Do not enter a value for the client secret.
- e. Select **Authorize** to begin a session.
- f. Close the window.

3. Close the **Available authorizations** dialog box.

4. Select **GET/assets**.

5. Select **Try it out**.

6. Select **Execute**.

7. Scroll down in the response body to the **Compute** section, and copy the `parent` and `id` values for the failed compute node.

8. Select **DELETE/assets/{asset_id}/compute-nodes/{compute_id}**.

9. Select **Try it out**.

10. Enter the `parent` and `id` values you copied in a previous step.

11. Select **Execute**.

Configure a proxy server

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

A proxy server is used for telemetry collectors and reverse tunnel connections. You can enable and configure a proxy server using the REST API UI if you did not already configure a proxy server during installation or upgrade. You can also modify existing proxy server settings or disable a proxy server.

The command to configure a proxy server updates and then returns the current proxy settings for the management node. The proxy settings are used by Active IQ, the NetApp HCI monitoring service that is deployed by the NetApp Deployment Engine, and other Element software utilities that are installed on the management node, including the reverse support tunnel for NetApp Support.

What you'll need

- You should know host and credential information for the proxy server you are configuring.
- Ensure that your cluster version is running NetApp Element software 11.3 or later.
- Ensure that you have deployed a management node running version 11.3 or later.
- (Management node 12.0 and 12.2) You have updated NetApp Hybrid Cloud Control to management services version 2.16 before configuring a proxy server.

Steps

1. Access the REST API UI on the management node by entering the management node IP address followed by /mnode:

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

2. Select **Authorize** or any lock icon and complete the following:

- a. Enter the cluster user name and password.
- b. Enter the client ID as mnode-client.
- c. Select **Authorize** to begin a session.
- d. Close the window.

3. Select **PUT /settings**.

4. Select **Try it out**.

5. To enable a proxy server, you must set `use_proxy` to true. Enter the IP or host name and proxy port destinations.

The proxy user name, proxy password, and SSH port are optional and should be omitted if not used.

```
{
  "proxy_ip_or_hostname": "[IP or name]",
  "use_proxy": [true/false],
  "proxy_username": "[username]",
  "proxy_password": "[password]",
  "proxy_port": [port value],
  "proxy_ssh_port": [port value: default is 443]
}
```

6. Select **Execute**.



You might need to reboot your management node depending on your environment.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Verify management node OS and services versions

You can verify the version numbers of the management node OS, management services bundle, and individual services running on the management node using the REST API in the management node.

What you'll need

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

Options

- [API commands](#)
- [REST API UI steps](#)

API commands

- Get version information about the management node OS, the management services bundle, and the management node API (mnode-api) service that are running on the management node:

```
curl -X GET "https://<ManagementNodeIP>/mnode/about" -H "accept: application/json"
```

- Get version information about individual services running on the management node:

```
curl -X GET "https://<ManagementNodeIP>/mnode/services?status=running" -H "accept: */*" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer \${TOKEN} used by the API command when you [authorize](#). The bearer \${TOKEN} is in the curl response.

REST API UI steps

1. Access the REST API UI for the service by entering the management node IP address followed by /mnode/:

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

2. Do one of the following:

- Get version information about the management node OS, the management services bundle, and the management node API (mnode-api) service that are running on the management node:
 - a. Select **GET /about**.
 - b. Select **Try it out**.
 - c. Select **Execute**.

The management services bundle version ("mnode_bundle_version"), management node OS version ("os_version"), and management node API version ("version") are indicated in the response body.

- Get version information about individual services running on the management node:
 - a. Select **GET /services**.
 - b. Select **Try it out**.
 - c. Select the status as **Running**.
 - d. Select **Execute**.

The services that are running on the management node are indicated in the response body.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Getting logs from management services

You can retrieve logs from the services running on the management node using the REST API. You can pull logs from all public services or specify specific services and use query parameters to better define the return results.

What you'll need

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

Steps

1. Open the REST API UI on the management node.

- Beginning with management services 2.21.61:

```
https://<ManagementNodeIP>/mnode/4/
```

- For management services 2.20.69 or earlier:

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

2. Select **Authorize** or any lock icon and complete the following:

- Enter the cluster user name and password.
- Enter the client ID as mnode-client if the value is not already populated.
- Select **Authorize** to begin a session.
- Close the window.

3. Select **GET /logs**.

4. Select **Try it out**.

5. Specify the following parameters:

- **Lines**: Enter the number of lines you want the log to return. This parameter is an integer that defaults to 1000.



Avoid requesting the entire history of log content by setting Lines to 0.

- **since**: Adds a ISO-8601 timestamp for the service logs starting point.



Use a reasonable `since` parameter when gathering logs of wider timespans.

- service-name: Enter a service name.



Use the GET /services command to list services on the management node.

- stopped: Set to true to retrieve logs from stopped services.

6. Select **Execute**.

7. From the response body, select **Download** to save the log output.

Find more Information

- [NetApp Element Plug-in for vCenter Server](#)

Manage support connections

Start a remote NetApp Support session

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

About this task

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

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



By default, the capability for remote access is enabled on the management node. To disable remote access functionality, see [Manage SSH functionality on the management node](#). You can enable remote access functionality again, if needed.

Steps

- Log in to your management node and open a terminal session.
- At a prompt, enter the following:

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

- To close the remote support tunnel, enter the following:

```
rst --killall
```

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Manage SSH functionality on the management node

You can disable, re-enable, or determine the status of the SSH capability on the management node (mNode) using the REST API. SSH capability that provides [NetApp Support remote support tunnel \(RST\) session access](#) is enabled on the management node by default.

Beginning with Management Services 2.20.69, you can enable and disable SSH capability on the management node using the NetApp Hybrid Cloud Control UI.

What you'll need

- **NetApp Hybrid Cloud Control permissions:** You have permissions as administrator.
- **Cluster administrator permissions:** You have permissions as administrator on the storage cluster.
- **Element software:** Your cluster is running NetApp Element software 11.3 or later.
- **Management node:** You have deployed a management node running version 11.3 or later.
- **Management services updates:**
 - To use the NetApp Hybrid Cloud Control UI, you have updated your [management services bundle](#) to version 2.20.69 or later.
 - To use the REST API UI, you have updated your [management services bundle](#) to version 2.17.

Options

- [Disable or enable the SSH capability on the management node using NetApp Hybrid Cloud Control UI](#)

You can do any of the following tasks after you [authenticate](#):

- [Disable or enable the SSH capability on the management node using APIs](#)
- [Determine status of the SSH capability on the management node using APIs](#)

Disable or enable the SSH capability on the management node using NetApp Hybrid Cloud Control UI

You can disable or re-enable SSH capability on the management node. SSH capability that provides [NetApp Support remote support tunnel \(RST\) session access](#) is disabled by default on management nodes running management services 2.18 or later. Disabling SSH does not terminate or disconnect existing SSH client sessions to the management node. If you disable SSH and elect to re-enable it at a later time, you can do so using the NetApp Hybrid Cloud Control UI.



To enable or disable support access using SSH for a storage cluster, you must use the [Element UI cluster settings page](#).

Steps

1. From the Dashboard, select the options menu on the top right and select **Configure**.
2. In the **Support Access for Management Node** screen, toggle the switch to enable management node SSH.
3. After you complete troubleshooting, in the **Support Access for Management Node** screen, toggle the switch to disable management node SSH.

Disable or enable the SSH capability on the management node using APIs

You can disable or re-enable SSH capability on the management node. SSH capability that provides [NetApp Support remote support tunnel \(RST\) session access](#) is enabled on the management node by default. Disabling SSH does not terminate or disconnect existing SSH client sessions to the management node. If you disable SSH and elect to re-enable it at a later time, you can do so using the same API.

API command

For management services 2.18 or later:

```
curl -k -X PUT
"https://<>ManagementNodeIP>/mnode/2/settings/ssh?enabled=<false/true>" -H
"accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

For management services 2.17 or earlier:

```
curl -X PUT
"https://<ManagementNodeIP>/mnode/settings/ssh?enabled=<false/true>" -H
"accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer \${TOKEN} used by the API command when you [authorize](#). The bearer \${TOKEN} is in the curl response.

REST API UI steps

1. Access the REST API UI for the management node API service by entering the management node IP address followed by /mnode/:

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

2. Select **Authorize** and complete the following:

- a. Enter the cluster user name and password.
- b. Enter the client ID as mnode-client.
- c. Select **Authorize** to begin a session.
- d. Close the window.

3. From the REST API UI, select **PUT /settings/ssh**.

- a. Select **Try it out**.
- b. Set the **enabled** parameter to **false** to disable SSH or **true** to re-enable SSH capability that you previously disabled.
- c. Select **Execute**.

Determine status of the SSH capability on the management node using APIs

You can determine whether or not SSH capability is enabled on the management node using a management node service API. SSH is enabled by default on the management node.

API command

For management services 2.18 or later:

```
curl -k -X PUT  
"https://<<ManagementNodeIP>/mnode/2/settings/ssh?enabled=<false/true>" -H  
"accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

For management services 2.17 or earlier:

```
curl -X PUT  
"https://<ManagementNodeIP>/mnode/settings/ssh?enabled=<false/true>" -H  
"accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer \${TOKEN} used by the API command when you **authorize**. The bearer \${TOKEN} is in the curl response.

REST API UI steps

1. Access the REST API UI for the management node API service by entering the management node IP address followed by /mnode/:

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

2. Select **Authorize** and complete the following:

- a. Enter the cluster user name and password.
- b. Enter the client ID as mnode-client.
- c. Select **Authorize** to begin a session.
- d. Close the window.

3. From the REST API UI, select **GET /settings/ssh**.

- a. Select **Try it out**.
- b. Select **Execute**.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Power your NetApp HCI system off or on

Powering your NetApp HCI system off or on

You can power off or power on your NetApp HCI system if you have a scheduled outage, need to perform hardware maintenance, or need to expand the system. Use the following tasks to power off or power on your NetApp HCI system as required.

You might need to power off your NetApp HCI system under a number of different circumstances, such as:

- Scheduled outages
- Chassis fan replacements
- Firmware upgrades
- Storage or compute resource expansion

The following is an overview of the tasks you need to complete to power off a NetApp HCI system:

- Power off all virtual machines except the VMware vCenter server (vCSA).
- Power off all ESXi servers except the one hosting the vCSA.
- Power off the vCSA.
- Power off the NetApp HCI storage system.

The following is an overview of the tasks you need to complete to power on a NetApp HCI system:

- Power on all physical storage nodes.
- Power on all physical compute nodes.
- Power on the vCSA.
- Verify the system and power on additional virtual machines.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)

Power off compute resources for a NetApp HCI system

To power off NetApp HCI compute resources, you need to power off individual VMware ESXi hosts as well as the VMware vCenter Server Appliance in a certain order.

Steps

1. Log in to the vCenter instance controlling the NetApp HCI system and determine the ESXi machine hosting the vCenter Server Virtual Appliance (vCSA).
2. After you have determined the ESXi host running the vCSA, power down all other virtual machines other than the vCSA as follows:
 - a. Select a virtual machine.
 - b. Right-click and select **Power > Shut Down Guest OS**.
3. Power off all ESXi hosts that are not the ESXi host running the vCSA.
4. Power off the vCSA.

This will cause the vCenter session to end because the vCSA disconnects during the power-off process. All virtual machines should now be shut down with only one ESXi host powered on.

5. Log in to the running ESXi host.
6. Verify that all virtual machines on the host are powered off.

7. Shut down the ESXi host.

This disconnects any iSCSI sessions open to the NetApp HCI storage cluster.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)

Power off storage resources for a NetApp HCI system

When you power off storage resources for NetApp HCI, you need to use the Shutdown Element API method to properly halt the storage nodes.

Steps

After you power off the compute resources, you use a web browser to shut down all the nodes of the NetApp HCI storage cluster.

1. Log in to the storage cluster and verify that you are connected to the correct MVIP.
2. (Optional) Verify that all I/O operations from the hosts have stopped:
 - a. Quiesce I/O from the host side by using the appropriate commands for the one or more hypervisors in use.
 - b. In the Cluster UI, select **Reporting > Overview**. There should be no activity on the "Cluster Input/Output" graph.
 - c. After all I/O operations have stopped, wait for 20 minutes before shutting down the cluster.
3. Verify that the iSCSI session count is zero.
4. Navigate to **Cluster > Nodes > Active**, and record the node IDs for all of the active nodes in the cluster.
5. To power off the NetApp HCI storage cluster, open a web browser and use the following URL to invoke the power off and halt procedure, where `{MVIP}` is the management IP address of the NetApp HCI storage system and the `nodes=[]` array includes the node IDs that you recorded in step 4. For example:

```
https://{}{MVIP}/json-rpc/1.0?method=Shutdown&nodes=[1,2,3,4]&option=halt
```



You can run the command in an incognito window to avoid running it again at a later stage from the saved URL.

6. Enter the cluster administrator user name and password.
7. Validate that the API call returned successfully by verifying that all storage cluster nodes are included in the `successful` section of the API result.

You have successfully powered off all the NetApp HCI storage nodes.

8. Close the browser or tab to avoid selecting the "back" button and repeating the API call.

When you restart the cluster, you must follow certain steps to verify that all nodes come online:

1. Verify that all Critical severity and `volumesOffline` cluster faults have been resolved.
2. Wait for 10 to 15 minutes for the cluster to settle.
3. Start bringing up the hosts to access the data.

 If you want to allow more time when powering on nodes and verifying that they are healthy after maintenance, contact technical support for assistance with delaying data synchronization to prevent unnecessary bin syncing.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI](#) and [firmware versions for NetApp HCI storage nodes](#)

Power on storage resources for a NetApp HCI system

You can power on NetApp HCI after the scheduled outage is complete.

Steps

1. Power on all the storage nodes using either the physical power button or the BMC.
2. If using the BMC, log in to each node and navigate to **Remote Control > Power Control > Power On Server**.
3. When all the storage nodes are online, log in to the NetApp HCI storage system and verify that all nodes are operational.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI](#) and [firmware versions for NetApp HCI storage nodes](#)

Power on compute resources for a NetApp HCI system

You can power on compute resources for a NetApp HCI system after the scheduled outage is complete.

Steps

1. Power on compute nodes using the same steps you performed for powering on the storage nodes.
2. When all the compute nodes are operational, log in to the ESXi host that was running the vCSA.
3. Log in to the compute host and verify that it sees all the NetApp HCI datastores. For a typical NetApp HCI system, you should see all the ESXi local datastores and at least the following shared datastores:

NetApp-HCI-Datastore-[01,02]

1. Assuming all storage is accessible, power on the vCSA and any other required virtual machines as follows:
 - a. Select the virtual machines in the navigator, select all the virtual machines that you want to power on, then select the **Power on** button.

2. After you power on the virtual machines, wait for approximately 5 minutes and then use a web browser to navigate to the IP address or FQDN of the vCSA appliance.

If you do not wait long enough, a message appears stating that the vSphere Client web server is initializing.

3. After the vSphere Client initializes, log in and verify that all ESXi hosts and virtual machines are online.

Find more information

- [Supported firmware and ESXi driver versions for NetApp HCI and firmware versions for NetApp HCI storage nodes](#)

Monitor your NetApp HCI system with NetApp Hybrid Cloud Control

Monitor storage and compute resources on the Hybrid Cloud Control Dashboard

With the NetApp Hybrid Cloud Control Dashboard, you can view all your storage and compute resources at a glance. Additionally, you can monitor storage capacity, storage performance, and compute utilization.



When you launch a new NetApp Hybrid Cloud Control session for the first time, there might be a delay with loading the NetApp Hybrid Cloud Control Dashboard view when the management node is managing many clusters. The loading time varies depending on the number of clusters being actively managed by the management node. For subsequent launches, you will experience faster loading times.

Only compute nodes that are managed and clusters with at least one managed node in H-series hardware appear on the Hybrid Cloud Control Dashboard.

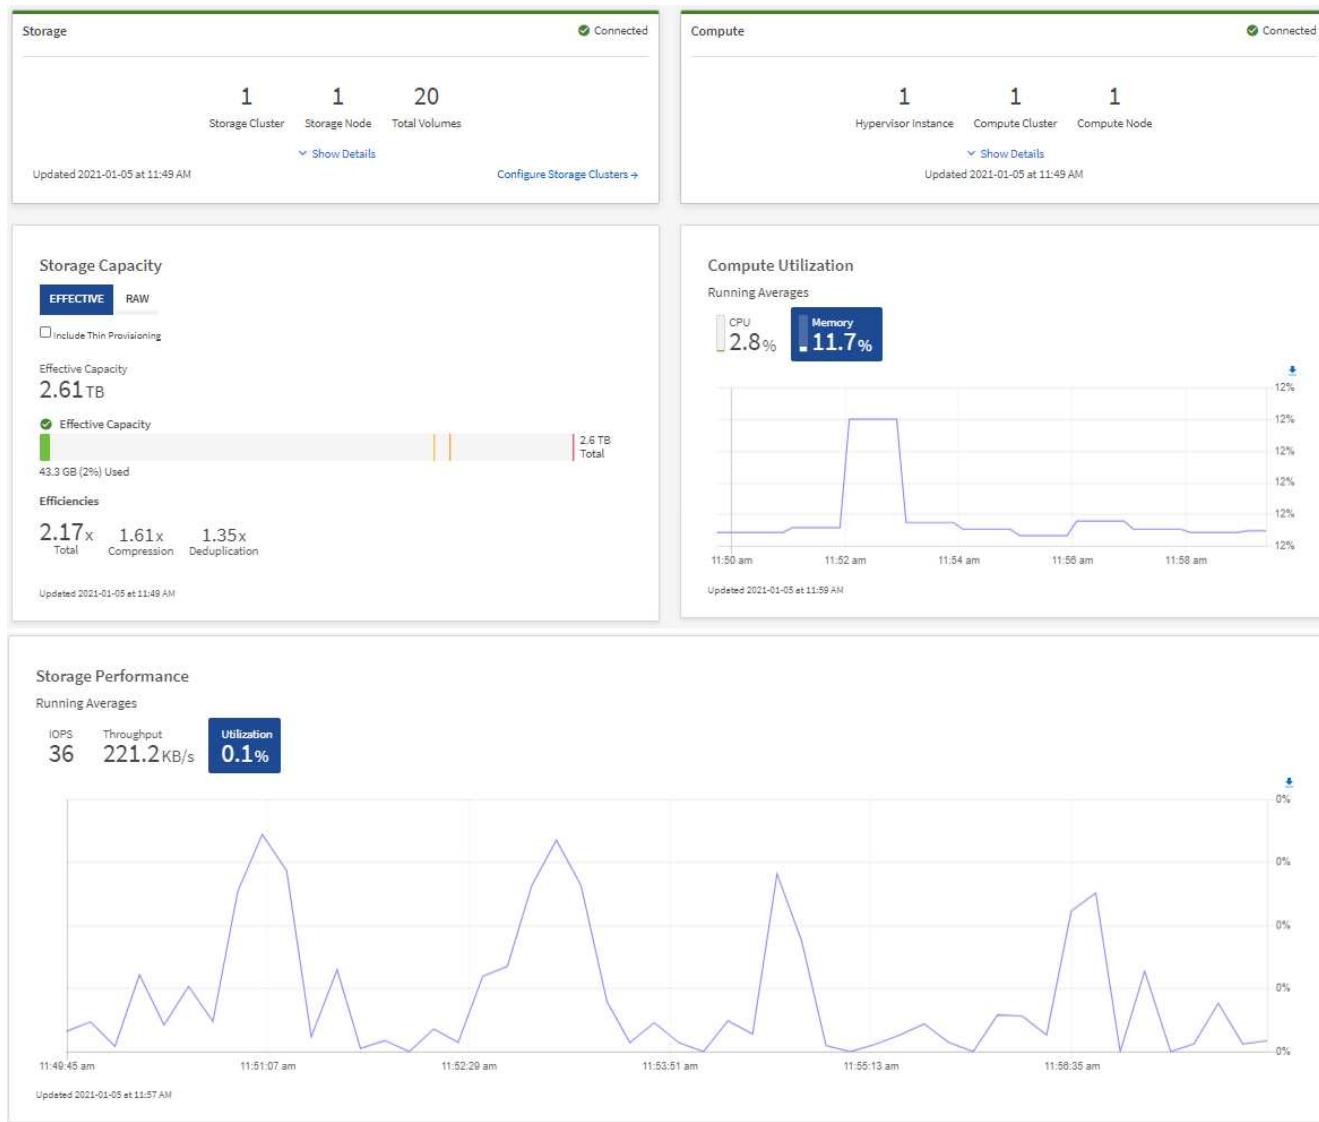
- [Access the NetApp HCC Dashboard](#)
- [Monitor storage resources](#)
- [Monitor compute resources](#)
- [Monitor storage capacity](#)
- [Monitor storage performance](#)
- [Monitor compute utilization](#)

Access the NetApp HCC Dashboard

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.
3. View the Hybrid Cloud Control Dashboard.



You might see some or all these panes, depending on your installation. For example, for storage-only installations, the Hybrid Cloud Control Dashboard shows only the Storage pane, the Storage Capacity pane, and the Storage Performance pane.

Monitor storage resources

Use the **Storage** pane to see your total storage environment. You can monitor the number of storage clusters, storage nodes, and total volumes.

To see details, in the Storage pane, select **Show Details**.

Cluster Name ↑	Nodes	Volumes	Connection Status
hci-tt-test8-cluster	4	16	Connected

i The Total Storage Nodes number does not include Witness Nodes from two-node storage clusters. The Witness Nodes are included in the Nodes number in the details section for that cluster.

i To see the most recent storage cluster data, use the Storage Clusters page, where polling occurs more frequently than on the Dashboard.

Monitor compute resources

Use the **Compute** pane to see your total NetApp H-series compute environment. You can monitor the number of compute clusters and total compute nodes.

To see details, in the Compute panes, select **Show Details**.

i Your vCenter instances only show in the Compute pane when at least one NetApp HCI compute node is associated with that instance. To list the vCenter instances linked in NetApp Hybrid Cloud Control, you can use the [APIs](#).

i To manage a compute node in NetApp Hybrid Cloud Control, you must [add the compute node to a vCenter host cluster](#).

Monitor storage capacity

Monitoring the storage capacity of your environment is critical. Using the Storage Capacity pane, you can determine your storage capacity efficiency gains with or without compression, deduplication, and thin provisioning features enabled.

You can see the total physical storage space available in your cluster on the **RAW** tab, and information about the provisioned storage on the **EFFECTIVE** tab.



To view cluster health, also look at the SolidFire Active IQ Dashboard. See [Monitor performance, capacity, and cluster health in NetApp SolidFire Active IQ](#).

Steps

1. Select the **RAW** tab, to see the total physical storage space used and available in your cluster.

Look at the vertical lines to determine whether your used capacity is less than the total or less than Warning, Error, or Critical thresholds. Hover over the lines to see details.



You can set the threshold for Warning, which defaults to 3% below the Error threshold. The Error and Critical thresholds are preset and not configurable by design. The Error threshold indicates that less than one node of capacity remains in the cluster. For steps on setting the threshold, see [Setting cluster full threshold](#).



For details about the related cluster thresholds Element API, see ["getClusterFullThreshold"](#) in the *Element API Guide*. To view details about block and metadata capacity, see [Understanding cluster fullness levels](#) in the *Element User Guide*.

2. Select the **EFFECTIVE** tab, to see information about total storage provisioned to connected hosts and to see efficiency ratings.
 - a. Optionally, check **Include Thin Provisioning** to see thin provisioning efficiency rates in the Effective Capacity bar chart.
 - b. **Effective Capacity bar chart:** Look at the vertical lines to determine whether your used capacity is less than the total or less than Warning, Error, or Critical thresholds. Similar to the Raw tab, you can hover over the vertical lines to see details.
 - c. **Efficiencies:** Look at these ratings to determine your storage capacity efficiency gains with compression, deduplication, and thin provisioning features enabled. For example, if compression shows as "1.3x", this means that storage efficiency with compression enabled is 1.3 times more efficient than without it.



Total Efficiencies equals $(\text{maxUsedSpace} * \text{efficiency factor}) / 2$, where $\text{efficiencyFactor} = (\text{thinProvisioningFactor} * \text{deDuplicationFactor} * \text{compressionFactor})$. When Thin Provisioning is unchecked, it is not included in the Total Efficiency.

- d. If the effective storage capacity nears an Error or Critical threshold, consider clearing the data on your system. Alternatively, consider expanding your system.

See [Expansion overview](#).

3. For further analysis and historical context, look at [NetApp SolidFire Active IQ details](#).

Monitor storage performance

You can look at how much IOPS or throughput you can get out of a cluster without surpassing the useful performance of that resource by using the Storage Performance pane. Storage performance is the point at which you get the maximum utilization before latency becomes an issue.

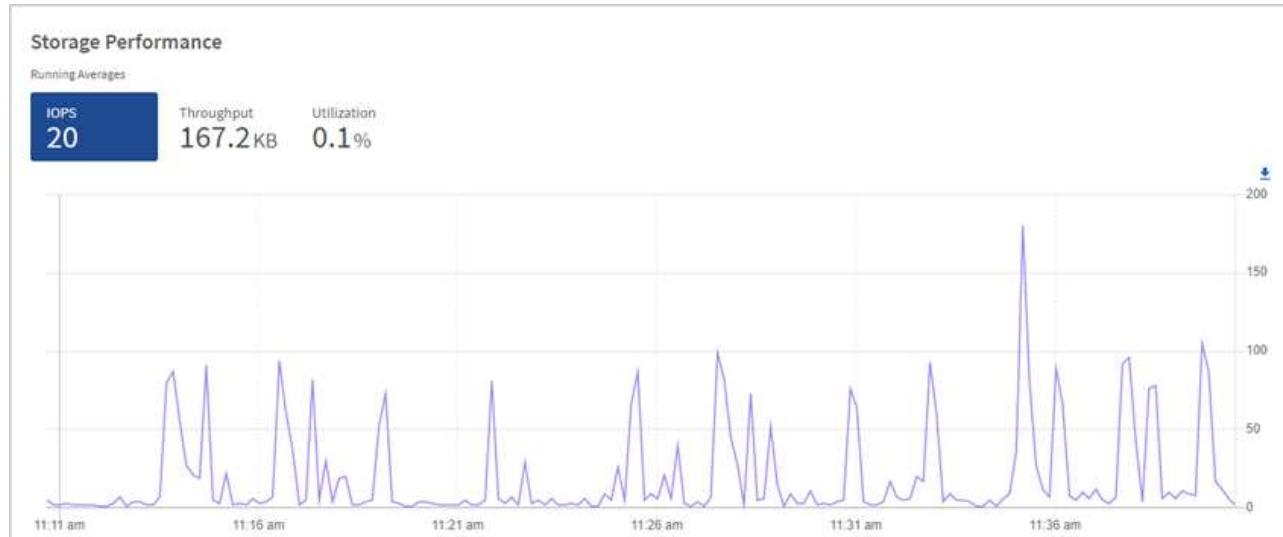
The Storage Performance pane helps you identify whether the performance is reaching the point where the performance might degrade if the workloads increase.

The information on this pane refreshes every 10 seconds and shows an average of all the points on the graph.

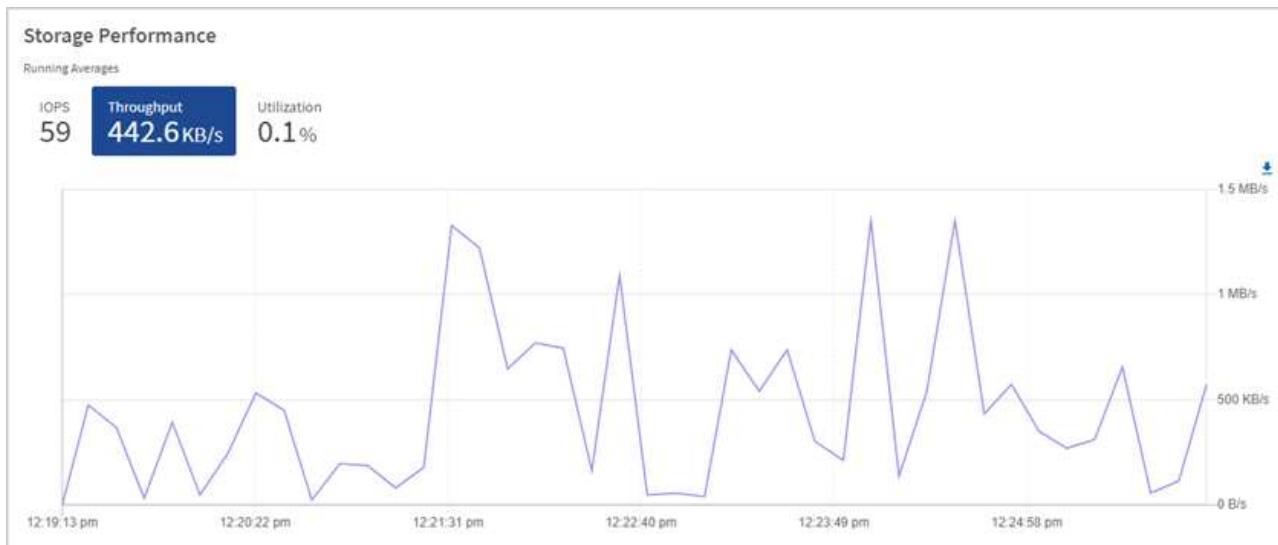
For details about the associated Element API method, see the [GetClusterStats](#) method in the *Element API Reference Guide*.

Steps

1. View the Storage Performance pane. For details, hover over points in the graph.
 - a. **IOPS** tab: See the current operations per second. Look for trends in data or spikes. For example, if you see that the maximum IOPS is 160K and 100K of that is free or available IOPS, you might consider adding more workloads to this cluster. On the other hand, if you see that only 140K is available, you might consider offloading workloads or expanding your system.



- b. **Throughput** tab: Monitor patterns or spikes in throughput. Also monitor for continuously high throughput values, which might indicate that you are nearing the maximum useful performance of the resource.



- c. **Utilization** tab: Monitor the utilization of IOPS in relation to the total IOPS available summed up at the cluster level.



2. For further analysis, look at storage performance by using the NetApp Element Plug-in for vCenter Server.

[Performance shown in the NetApp Element Plug-in for vCenter Server.](#)

Monitor compute utilization

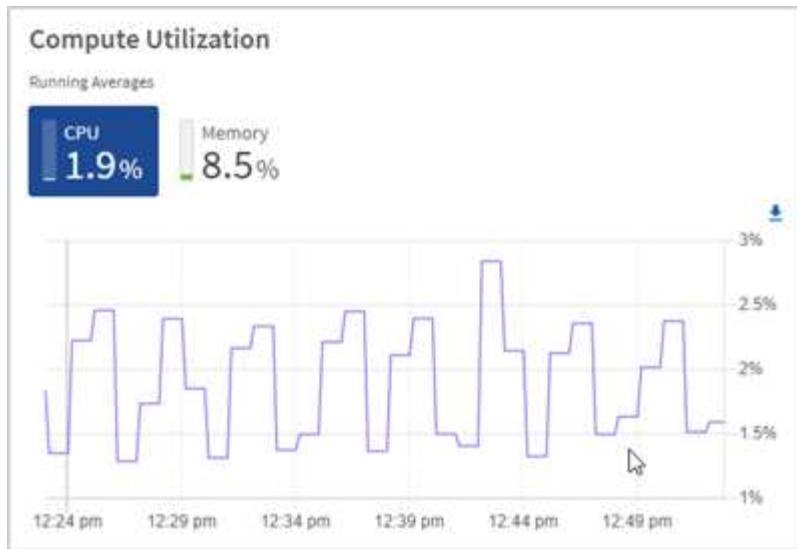
In addition to monitoring IOPS and throughput of your storage resources, you also might want to view the CPU and memory usage of your compute assets. The total IOPS that a node can provide is based on the physical characteristics of the node, for example, the number of CPUs, the CPU speed, and the amount of RAM.

Steps

1. View the **Compute Utilization** pane. Using both the CPU and Memory tabs, look for patterns or spikes in utilization. Also look for continuously high usage, indicating that you might be nearing the maximum utilization for the compute clusters.



This pane shows data only for those compute clusters managed by this installation.



- a. **CPU** tab: See the current average of CPU utilization on the compute cluster.
 - b. **Memory** tab: See the current average memory usage on the compute cluster.
2. For further analysis on compute information, see [NetApp SolidFire Active IQ for historical data](#).

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp SolidFire Active IQ Documentation](#)

View your inventory on the Nodes page

You can view both your storage and compute assets in your system and determine their IP addresses, names, and software versions.

You can view storage information for your multiple node systems and any NetApp HCI Witness Nodes associated with two-node or three-node clusters.

Witness Nodes manage quorum within the cluster; they are not used for storage. Witness Nodes are applicable only to NetApp HCI and not to all-flash storage environments.

For more information about Witness Nodes, see [Nodes definitions](#).

For SolidFire Enterprise SDS nodes, you can monitor inventory on the Storage tab.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.

The NetApp Hybrid Cloud Control Dashboard appears.

3. In the left navigation, select **Nodes**.

Nodes

The screenshot shows the 'Nodes' page in the NetApp HCI interface. At the top, a yellow banner states: 'Only NetApp HCI Nodes are displayed on this page.' Below this, the 'STORAGE' tab is selected, showing 'Cluster 1' with '1 of 1 Two-node'. The main table lists two storage nodes: 'stg01' and 'stg02', both with 'H410S-0' Node Model, '12.0.0.318' Element Version, and 'Management IP Address' ending in '- VLAN 1184'. Below the table, it says '1-2 of 2 results' with a page navigation bar. A 'Witness Nodes' section follows, showing a table with two rows: 'wit01' and 'wit02', each with a 'Management IP Address' and a 'Storage (iSCSI) IP Address'.



When you launch a new NetApp Hybrid Cloud Control session for the first time, there might be a delay with loading the NetApp Hybrid Cloud Control Nodes page when the management node is managing many clusters. The loading time varies depending on the number of clusters being actively managed by the management node. For subsequent launches, you will experience faster loading times.

4. On the **Storage** tab of the Nodes page, review the following information:
- Two-node clusters: A “two-node” label appears on the Storage tab and the associated Witness Nodes are listed.
 - Three-node clusters: The storage nodes and associated Witness Nodes are listed. Three-node clusters have a Witness Node deployed on standby to maintain high availability in the case of node failure.
 - Clusters with four nodes or more: Information for clusters with four or more nodes appears. Witness Nodes do not apply. If you started with two or three storage nodes and added more nodes, the Witness Nodes table still appears. Otherwise, the Witness Nodes table does not appear.
 - The firmware bundle version: Starting with management services version 2.14, if you have clusters running Element 12.0 or later, you can see the firmware bundle version for these clusters. If the nodes in a cluster have different firmware versions on them, you can see **Multiple** in the **Firmware Bundle Version** column.
5. To view compute inventory information, select **Compute**.
6. You can manipulate the information on these pages in several ways:
- To filter the list of items in the results, select the **Filter** icon and select the filters. You can also enter text for the filter.
 - To show or hide columns, select the **Show/Hide Columns** icon.
 - To download the table, select the **Download** icon.
 - To add or edit the stored BMC credentials for a compute node with BMC connection errors, select **Edit connection settings** in the error message text in the **BMC Connection Status** column. Only if the

connection attempt fails for a compute node, an error message is displayed in this column for that node.



To view the number of storage and compute resources, look at the NetApp Hybrid Cloud Control (HCC) Dashboard. See [Monitor storage and compute resources with the HCC Dashboard](#).



To manage a compute node in NetApp Hybrid Cloud Control, you must [add the compute node to a vCenter host cluster](#).

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Edit Baseboard Management Controller connection information

You can change Baseboard Management Controller (BMC) administrator credentials in NetApp Hybrid Cloud Control for each of your compute nodes. You might need to change credentials prior to upgrading BMC firmware or to resolve a Hardware ID not available or Unable to Detect error indicated in NetApp Hybrid Cloud Control.

What you'll need

Cluster administrator permissions to change BMC credentials.



If you set BMC credentials during a health check, there can be a delay of up to 15 minutes before the change is reflected on the **Nodes** page.

Options

Choose one of the following options to change BMC credentials:

- [Use NetApp Hybrid Cloud Control to edit BMC information](#)
- [Use the REST API to edit BMC information](#)

Use NetApp Hybrid Cloud Control to edit BMC information

You can edit the stored BMC credentials using the NetApp Hybrid Cloud Control Dashboard.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.

3. In the left navigation blue box, select the NetApp HCI installation.

The NetApp Hybrid Cloud Control Dashboard appears.

4. In the left navigation, select **Nodes**.
5. To view compute inventory information, select **Compute**.

A list of your compute nodes appears. The **BMC Connection Status** column shows the result of BMC connection attempts for each compute node. If the connection attempt fails for a compute node, an error message is displayed in this column for that node.

6. To add or edit the stored BMC credentials for a compute node with BMC connection errors, select **Edit connection settings** in the error message text.
7. In the dialog that appears, add the correct administrator user name and password for the BMC of this compute node.
8. Select **Save**.
9. Repeat steps 6 through 8 for any compute node that has missing or incorrect stored BMC credentials.



Updating BMC information refreshes the inventory and ensures that management node services are aware of all hardware parameters needed to complete the upgrade.

Use the REST API to edit BMC information

You can edit the stored BMC credentials using the NetApp Hybrid Cloud Control REST API.

Steps

1. Locate the compute node hardware tag and BMC information:
 - a. Open the inventory service REST API UI on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```
 - b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client`.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the authorization window.
 - c. From the REST API UI, select **GET /installations**.
 - d. Select **Try it out**.
 - e. Select **Execute**.
 - f. From the response, copy the installation asset ID (`id`).
 - g. From the REST API UI, select **GET /installations/{id}**.
 - h. Select **Try it out**.
 - i. Paste the installation asset ID into the `id` field.
 - j. Select **Execute**.
 - k. From the response, copy and save the node asset id (`id`), BMC IP address (`bmcAddress`), and node serial number (`chassisSerialNumber`) for use in a later step.

```

"nodes": [
  {
    "bmcDetails": {
      "bmcAddress": "10.117.1.111",
      "credentialsAvailable": false,
      "credentialsValidated": false
    },
    "chassisSerialNumber": "221111019323",
    "chassisSlot": "C",
    "hardwareId": null,
    "hardwareTag": "00000000-0000-0000-ac1f6ab4ecf6",
    "id": "8cd91e3c-1b1e-1111-b00a-4c9c4900b000",
  }
]

```

2. Open the hardware service REST API UI on the management node:

```
https://<ManagementNodeIP>/hardware/2/
```

3. Select **Authorize** and complete the following:

- Enter the cluster user name and password.
- Enter the client ID as `mnode-client` if the value is not already populated.
- Select **Authorize** to begin a session.
- Close the window.

4. Select **PUT /nodes/{hardware_id}**.

5. Select **Try it out**.

6. Enter the node asset id that you saved earlier in the `hardware_id` parameter.

7. Enter the following information in the payload:

Parameter	Description
assetId	The installation asset id (<code>id</code>) that you saved in step 1(f).
bmcIp	The BMC IP address (<code>bmcAddress</code>) that you saved in step 1(k).
bmcPassword	An updated password to log into the BMC.
bmcUsername	An updated user name to log into the BMC.
serialNumber	The chassis serial number of the hardware.

Example payload:

```
{  
  "assetId": "7bb41e3c-2e9c-2151-b00a-8a9b49c0b0fe",  
  "bmcIp": "10.117.1.111",  
  "bmcPassword": "mypassword1",  
  "bmcUsername": "admin1",  
  "serialNumber": "221111019323"  
}
```

8. Select **Execute** to update BMC credentials.

A successful result returns a response similar to the following:

```
{  
  "credentialid": "33333333-cccc-3333-cccc-333333333333",  
  "host_name": "hci-host",  
  "id": "8cd91e3c-1b1e-1111-b00a-4c9c4900b000",  
  "ip": "1.1.1.1",  
  "parent": "abcd01y3-ab30-1ccc-11ee-11f123zx7d1b",  
  "type": "BMC"  
}
```

Find more information

- [Known issues and workarounds for compute node upgrades](#)
- [NetApp Element Plug-in for vCenter Server](#)

Monitor volumes on your storage cluster

The SolidFire system provisions storage using volumes. Volumes are block devices accessed over the network by iSCSI or Fibre Channel clients. You can monitor details about access groups, accounts, initiators, used capacity, Snapshot data protection status, number of iSCSI sessions, and the Quality of Service (QoS) policy associated with the volume.

You can also see details on active and deleted volumes.

With this view, you might first want to monitor the Used capacity column.

You can access this information only if you have NetApp Hybrid Cloud Control administrative privileges.

Steps

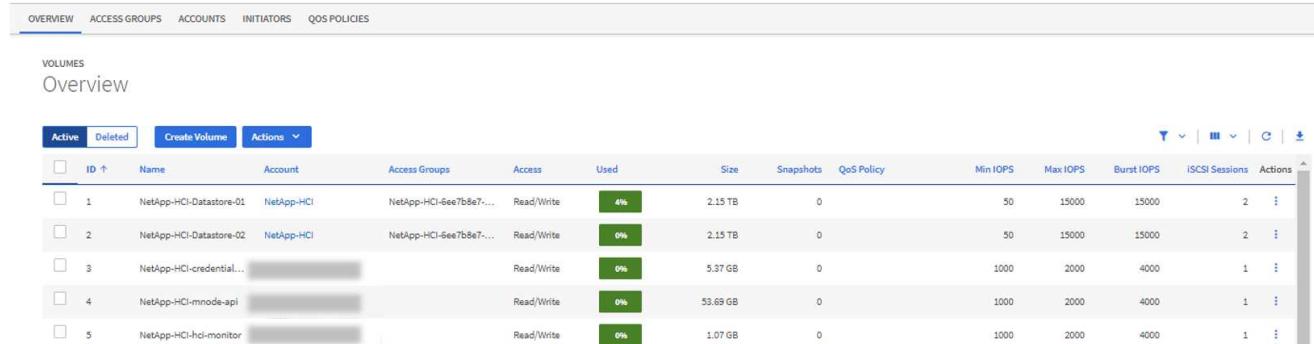
1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.
3. In the left navigation blue box, select the NetApp HCI installation.

The Hybrid Cloud Control Dashboard appears.

4. In the left navigation, select the cluster and select **Storage > Volumes**.



	ID	Name	Account	Access Groups	Access	Used	Size	Snapshots	QoS Policy	Min IOPS	Max IOPS	Burst IOPS	iSCSI Sessions	Actions
<input type="checkbox"/>	1	NetApp-HCI-Datastore-01	NetApp-HCI	NetApp-HCI-6ee7b8e7...	Read/Write	<div style="width: 4%;">4%</div>	2.15 TB	0		50	15000	15000	2	
<input type="checkbox"/>	2	NetApp-HCI-Datastore-02	NetApp-HCI	NetApp-HCI-6ee7b8e7...	Read/Write	<div style="width: 0%;">0%</div>	2.15 TB	0		50	15000	15000	2	
<input type="checkbox"/>	3	NetApp-HCI-credential...	[REDACTED]		Read/Write	<div style="width: 0%;">0%</div>	5.37 GB	0		1000	2000	4000	1	
<input type="checkbox"/>	4	NetApp-HCI-hnode-api	[REDACTED]		Read/Write	<div style="width: 0%;">0%</div>	53.69 GB	0		1000	2000	4000	1	
<input type="checkbox"/>	5	NetApp-HCI-hci-monitor	[REDACTED]		Read/Write	<div style="width: 0%;">0%</div>	1.07 GB	0		1000	2000	4000	1	

5. On the Volumes page, use the following options:



- a. Filter the results by selecting the **Filter** icon.
- b. Hide or show columns by selecting the **Hide/Show** icon.
- c. Refresh data by selecting the **Refresh** icon.
- d. Download a CSV file by selecting on the **Download** icon.
6. Monitor the Used capacity column. If Warning, Error, or Critical thresholds are reached, the color represents the used capacity status:
 - a. Warning - Yellow
 - b. Error - Orange
 - c. Critical - Red
7. From the Volumes view, select the tabs to see additional details about the volumes:
 - a. **Access Groups:** You can see the volume access groups that are mapped from initiators to a collection of volumes for secured access.
See information about [volume access groups](#).
 - b. **Accounts:** You can see the user accounts, which enable clients to connect to volumes on a node. When you create a volume, it is assigned to a specific user account.
See information about [NetApp HCI user accounts](#).
 - c. **Initiators:** You can see the iSCSI initiator IQN or Fibre Channel WWPNs for the volume. Each IQN added to an access group can access each volume in the group without requiring CHAP authentication. Each WWPN added to an access group enables Fibre Channel network access to the volumes in the access group.
See information about [access groups, initiators, and CHAP authentication methods](#) in the *NetApp*

- d. **QoS Policies:** You can see the QoS policy applied to the volume. A QoS policy applies standardized settings for minimum IOPS, maximum IOPS, and burst IOPS to multiple volumes.

See information about [performance and QoS policies](#).

See information about [Quality of Service policies](#) in the *NetApp Element User Guide*.

Find more information

- [NetApp SolidFire and Element Documentation Center](#)
- [NetApp Element Plug-in for vCenter Server](#)

Monitor performance, capacity, and cluster health with SolidFire Active IQ

By using SolidFire Active IQ, you can monitor the events, performance, and capacity of your clusters. You can access SolidFire Active IQ from the NetApp Hybrid Cloud Control Dashboard.

What you'll need

- You must have a NetApp Support account to take advantage of this service.
- You must have authorization to use management node REST APIs.
- You have deployed a management node running version 12.0 or later.
- Your cluster version is running NetApp Element software 12.0 or later.
- You have Internet access. The Active IQ collector service cannot be used from dark sites.

About this task

You can obtain continually updated historical views of cluster-wide statistics. You can set up notifications to alert you about specified events, thresholds, or metrics on a cluster so that they can be addressed quickly.

As part of your normal support contract, NetApp Support monitors this data and alerts you to potential system issues.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.
3. From the Dashboard, select the menu on the upper right.
4. Select **View Active IQ**.

The [SolidFire Active IQ Dashboard](#) appears.

5. To learn about SolidFire Active IQ, see the [SolidFire Active IQ Documentation](#).

You can also access the SolidFire Active IQ documentation from the dashboard by selecting the menu icon on the upper right and selecting **Documentation**.

6. From the SolidFire Active IQ interface, verify that the NetApp HCI compute and storage nodes are reporting telemetry correctly to Active IQ:

- a. If you have more than one NetApp HCI installation, select **Select a Cluster** and choose the cluster from the list.
- b. In the left navigation pane, select **Nodes**.

7. If a node or nodes are missing from the list, contact NetApp Support.



To view the number of storage and compute resources, look at the Hybrid Cloud Control (HCC) Dashboard. See [Monitor storage and compute resources with the HCC Dashboard](#).

Find more information

- [NetApp SolidFire Active IQ Documentation](#)
- [NetApp Element Plug-in for vCenter Server](#)

Collect logs for troubleshooting

If you have trouble with your NetApp HCI or SolidFire all-flash storage installation, you can collect logs to send to NetApp Support to help with diagnosis. You can either use NetApp Hybrid Cloud Control or the REST API to collect logs on NetApp HCI or Element systems.

What you'll need

- Ensure that your storage cluster version is running NetApp Element software 11.3 or later.
- Ensure that you have deployed a management node running version 11.3 or later.

Log collection options

Choose one of the following options:

- [Use NetApp Hybrid Cloud Control to collect logs](#)
- [Use the REST API to collect logs](#)

Use NetApp Hybrid Cloud Control to collect logs

You can access the log collection area from the NetApp Hybrid Cloud Control Dashboard.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster

administrator credentials.

3. From the Dashboard, select the menu on the upper right.
4. Select **Collect Logs**.

The **Collect Logs** page appears. If you have collected logs before, you can download the existing log package, or begin a new log collection.

5. Select a date range in the **Date Range** drop-down menu to specify what dates the logs should include.

If you specify a custom start date, you can select the date to begin the date range. Logs will be collected from that date up to the present time.

6. In the **Log Collection** section, select the types of log files the log package should include.

For storage and compute logs, you can expand the list of storage or compute nodes and select individual nodes to collect logs from (or all nodes in the list).

7. Select **Collect Logs** to start log collection.

Log collection runs in the background, and the page shows the progress.



Depending on the logs you collect, the progress bar might remain at a certain percentage for several minutes, or progress very slowly at some points.

8. Select **Download Logs** to download the log package.

The log package is in a compressed UNIX .tgz file format.

Use the REST API to collect logs

You can use REST API to collect NetApp HCI or Element logs.

Steps

1. Locate the storage cluster ID:
 - a. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/logs/1/
```
 - b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client` if the value is not already populated.
 - iii. Select **Authorize** to begin a session.
2. Collect logs from NetApp HCI or Element:
 - a. Select **POST /bundle**.
 - b. Select **Try it out**.
 - c. Change the values of the following parameters in the **Request body** field depending on which type of logs you need to collect and for what time range:

Parameter	Type	Description
modifiedSince	Date string	Only include logs modified after this date and time. For example, the value "2020-07-14T20:19:00.000Z" defines a start date of July 14, 2020 at 20:19 UTC.
computeLogs	Boolean	Set this parameter to <code>true</code> to include compute node logs.
computeIds	UUID array	If <code>computeLogs</code> is set to <code>true</code> , populate this parameter with the management node asset IDs of compute nodes to limit log collection to those specific compute nodes. Use the GET <a href="https://<ManagementNodeIP>/logs/1/bundle/options">https://<ManagementNodeIP>/logs/1/bundle/options endpoint to see all possible node IDs you can use.
mnodeLogs	Boolean	Set this parameter to <code>true</code> to include management node logs.
storageCrashDumps	Boolean	Set this parameter to <code>true</code> to include storage node crash debug logs.
storageLogs	Boolean	Set this parameter to <code>true</code> to include storage node logs.
storageNodeIds	UUID array	If <code>storageLogs</code> is set to <code>true</code> , populate this parameter with the storage cluster node IDs to limit log collection to those specific storage nodes. Use the GET <a href="https://<ManagementNodeIP>/logs/1/bundle/options">https://<ManagementNodeIP>/logs/1/bundle/options endpoint to see all possible node IDs you can use.

d. Select **Execute** to begin log collection.

The response should return a response similar to the following:

```
{
  "_links": {
    "self": "https://10.1.1.5/logs/1/bundle"
  },
  "taskId": "4157881b-z889-45ce-adb4-92b1843c53ee",
  "taskLink": "https://10.1.1.5/logs/1/bundle"
}
```

3. Check on the status of the log collection task:

- a. Select **GET /bundle**.
- b. Select **Try it out**.
- c. Select **Execute** to return a status of the collection task.
- d. Scroll to the bottom of the response body.

You should see a `percentComplete` attribute detailing the progress of the collection. If the collection is complete, the `downloadLink` attribute contains the full download link including the file name of the log package.

- e. Copy the file name at the end of the `downloadLink` attribute.

4. Download the collected log package:

- a. Select **GET /bundle/{filename}**.
- b. Select **Try it out**.
- c. Paste the file name you copied earlier into the `filename` parameter text field.
- d. Select **Execute**.

After execution, a download link appears in the response body area.

- e. Select **Download file** and save the resulting file to your computer.

The log package is in a compressed UNIX .tgz file format.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Upgrade your NetApp HCI system version 1.8

Upgrade sequence overview

You can keep your NetApp HCI system up-to-date after deployment by sequentially upgrading all NetApp HCI software components.

These components include management services, HealthTools, NetApp Hybrid Cloud Control, Element software, management node, compute firmware, compute drivers, and the Element Plug-in for vCenter Server.

 Beginning November 2023, you cannot start a component upgrade using NetApp Hybrid Cloud Control or REST API because the signing key certificates (private and public) expired on November 5, 2023. You can resolve this issue by following the workaround documented in the Knowledge Base article [SolidFire and HCI upgrades unable to start due to upgrade package upload error](#).

The [system upgrade sequence](#) content describes the tasks that are needed to complete a NetApp HCI system upgrade. Ideally you perform these procedures as part of the larger upgrade sequence and not in isolation. If a component-based upgrade or update is needed, see the procedure prerequisites to ensure additional complexities are addressed.

The [vSphere upgrade sequence](#) including Element Plug-in for vCenter Server content describes additional pre- and post-upgrade steps required to re-install the Element Plug-in for vCenter Server.

What you'll need

- You are running management node 11.3 or later. Newer versions of the management node have a modular architecture that provides individual services.



To check the version, log in to your management node and view the Element version number in the login banner. If you do not have 11.3, see [Upgrade your management node](#).

- You have upgraded your management services to at least version 2.1.326.

Upgrades using NetApp Hybrid Cloud Control are not available in earlier service bundle versions.

- You have ensured that the system time on all nodes is synced and that NTP is correctly configured for the storage cluster and nodes. Each node must be configured with a DNS nameserver in the per-node web UI (`https://[IP address]:442`) with no unresolved cluster faults related to time skew.

System upgrade sequence

Steps

1. [Update management services from Hybrid Cloud Control](#).



If you are updating management services to version 2.16 or later and you are running a management node 11.3 to 11.8, you will need to increase your management node VM's RAM prior to updating management services.



You must update to the latest management services bundle before upgrading your Element software.

2. (Optional) Upgrade to the latest HealthTools.



Upgrading HealthTools is only required if the management node and Element software you are running is 11.1 or earlier. HealthTools are not required for performing Element upgrades using NetApp Hybrid Cloud Control.

3. Run Element storage health checks prior to upgrading storage.

4. Upgrade your Element software and storage firmware.

5. (Optional) Upgrade your Element storage firmware only.



You might perform this task when a new storage firmware upgrade becomes available outside of a major release.

6. (Optional) Upgrade your management node.



Upgrading the management node operating system is no longer required to upgrade Element software on the storage cluster. If the management node is version 11.3 or higher, you can simply upgrade the management services to the latest version to perform Element upgrades using NetApp Hybrid Cloud Control. Follow the management node upgrade procedure for your scenario if you would like to upgrade the management node operating system for other reasons, such as security remediation.

7. Upgrade your Element Plug-in for vCenter Server.

8. Run compute node health checks prior to upgrading compute firmware.

9. Update your compute node drivers.

10. Update your compute node firmware using NetApp Hybrid Cloud Control or Automate your compute firmware upgrades with Ansible.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [Upgrade a NetApp SolidFire all-flash storage system](#)

System upgrade procedures

Update management services

You can update your management services to the latest bundle version after you have installed management node 11.3 or later.

Beginning with the Element 11.3 management node release, the management node design has been changed based on a new modular architecture that provides individual services. These modular services provide central and extended management functionality for NetApp HCI and SolidFire all-flash storage systems. Management services include system telemetry, logging, and update services, the QoSSIOC service for Element Plug-in for vCenter Server, NetApp Hybrid Cloud Control, and more.

About this task

- You must upgrade to the latest management services bundle before upgrading your Element software.



For the latest management services release notes describing major services, new features, bug fixes, and workarounds for each service bundle, see [the management services release notes](#).

What you'll need

Beginning with management services 2.20.69, you must accept and save the End User License Agreement (EULA) before using the NetApp Hybrid Cloud Control UI or API to upgrade management services:

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Update options

You can update management services using the NetApp Hybrid Cloud Control UI or the management node REST API:

- [Update management services using Hybrid Cloud Control](#) (Recommended method)
- [Update management services using the management node API](#)

Update management services using Hybrid Cloud Control

You can update your NetApp management services using NetApp Hybrid Cloud Control.

Management service bundles provide enhanced functionality and fixes to your installation outside of major releases.

What you'll need

- You are running management node 11.3 or later.
- If you are updating management services to version 2.16 or later and you are running a management node 11.3 to 11.8, you will need to increase your management node VM's RAM prior to updating management services:
 - a. Power off the management node VM.
 - b. Change the RAM of the management node VM from 12GB to 24GB RAM.
 - c. Power on the management node VM.
- Your cluster version is running NetApp Element software 11.3 or later.
- You have upgraded your management services to at least version 2.1.326. NetApp Hybrid Cloud Control upgrades are not available in earlier service bundles.



For a list of available services for each service bundle version, see the [Management Services Release Notes](#).

Steps

1. Open the IP address of the management node in a web browser:

`https://<ManagementNodeIP>`

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **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.



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

5. If your installation can access the internet and if a management services upgrade is available, select **Begin Upgrade**.
6. If your installation cannot access the internet, do the following:
 - a. Follow the instructions on the page to download and save a management services upgrade package to your computer.
 - b. Select **Browse** to locate the package you saved and upload it.

After you upload the package, the upgrade starts automatically.

After the upgrade begins, you can see the upgrade status on this page. During the upgrade, you might lose connection with NetApp Hybrid Cloud Control and have to log back in to see the results of the upgrade.

Update management services using the management node API

Users should ideally perform management services updates from NetApp Hybrid Cloud Control. You can however manually upload, extract, and deploy a service bundle update for management services to the management node using the REST API. You can run each command from the REST API UI for the management node.

What you'll need

- You have deployed a NetApp Element software management node 11.3 or later.
- If you are updating management services to version 2.16 or later and you are running a management node 11.3 to 11.8, you will need to increase your management node VM's RAM prior to updating management services:
 - a. Power off the management node VM.
 - b. Change the RAM of the management node VM from 12GB to 24GB RAM.
 - c. Power on the management node VM.
- Your cluster version is running NetApp Element software 11.3 or later.
- You have upgraded your management services to at least version 2.1.326. NetApp Hybrid Cloud Control upgrades are not available in earlier service bundles.



For a list of available services for each service bundle version, see the [Management Services Release Notes](#).

Steps

1. Open the REST API UI on the management node: <https://<ManagementNodeIP>/mnode>
2. Select **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. Select **Authorize** to begin a session.
 - d. Close the window.
3. Upload and extract the service bundle on the management node using this command: `PUT /services/upload`
4. Deploy the management services on the management node: `PUT /services/deploy`
5. Monitor the status of the update: `GET /services/update/status`

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

```
{  
  "current_version": "2.10.29",  
  "details": "Updated to version 2.17.52",  
  "status": "success"  
}
```

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Upgrade to the latest HealthTools

Before you begin the Element storage upgrade, you should upgrade your HealthTools suite. Upgrading HealthTools is only required if the management node and Element software you are running is 11.1 or earlier. HealthTools are not required for performing Element upgrades using NetApp Hybrid Cloud Control.

What you'll need

- You are running management node 11.0, 11.1 or later.
- You have upgraded your management services to at least version 2.1.326.

NetApp Hybrid Cloud Control upgrades are not available in earlier service bundle versions.

- You have downloaded the latest version of [HealthTools](#) and copied the installation file to the management node.



You can check the locally installed version of HealthTools by running the `sfupdate-healthtools -v` command.

- To use HealthTools with dark sites, you need to do these additional steps:

- Download a [JSON file](#) from the NetApp Support Site on a computer that is not the management node and rename it to `metadata.json`.
- Have the management node up and running at the dark site.

About this task

The commands in the HealthTools suite require escalated privileges to run. Either preface commands with `sudo` or escalate your user to root privileges.



The HealthTools version you use might be more up to date than the sample input and response below.

Steps

1. Run the `sfupdate-healthtools <path to install file>` command to install the new HealthTools software.

Sample input:

```
sfupdate-healthtools /tmp/solidfire-healthtools-2020.03.01.09.tgz
```

Sample response:

```
Checking key signature for file /tmp/solidfirehealthtools-2020.03.01.09/components.tgz
installing command sfupdate-healthtools
Restarting on version 2020.03.01.09
sfupdate-healthtools /sf/bin/sfupdate-healthtools -r 2020.03.01.09
installing command sfupgradecheck
installing command sfinstall
installing command sfresetupgrade
```

2. Run the `sfupdate-healthtools -v` command to verify the installed version has been upgraded.

Sample response:

```
Currently installed version of HealthTools:
2020.03.01.09
```

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Run Element storage health checks prior to upgrading storage

You must run health checks prior to upgrading Element storage to ensure all storage nodes in your cluster are ready for the next Element storage upgrade.

What you'll need

- **Management services:** You have updated to the latest management services bundle (2.10.27 or later).



You must upgrade to the latest management services bundle before upgrading your Element software.

- **Management node:** You are running management node 11.3 or later.
- **Element software:** Your cluster version is running NetApp Element software 11.3 or later.
- **End User License Agreement (EULA):** Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to run Element storage health checks:

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Health check options

You can run health checks using NetApp Hybrid Cloud Control (HCC) UI, HCC API, or the HealthTools suite:

- [Use NetApp Hybrid Cloud Control to run Element storage health checks prior to upgrading storage](#) (Preferred method)
- [Use API to run Element storage health checks prior to upgrading storage](#)
- [Use HealthTools to run Element storage health checks prior to upgrading storage](#)

You can also find out more about storage health checks that are run by the service:

- [Storage health checks made by the service](#)

Use NetApp Hybrid Cloud Control to run Element storage health checks prior to upgrading storage

Using NetApp Hybrid Cloud Control (HCC), you can verify that a storage cluster is ready to be upgraded.

Steps

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select the **Storage** tab.
5. Select the health check  for the cluster you want to check for upgrade readiness.

6. On the **Storage Health Check** page, select **Run Health Check**.
7. If there are issues, do the following:
 - a. Go to the specific KB article listed for each issue or perform the specified remedy.
 - b. If a KB is specified, complete the process described in the relevant KB article.
 - c. After you have resolved cluster issues, select **Re-Run Health Check**.

After the health check completes without errors, the storage cluster is ready to upgrade. See storage node upgrade [instructions](#) to proceed.

Use API to run Element storage health checks prior to upgrading storage

You can use REST API to verify that a storage cluster is ready to be upgraded. The health check verifies that there are no obstacles to upgrading, such as pending nodes, disk space issues, and cluster faults.

Steps

1. Locate the storage cluster ID:
 - a. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/mnode
```
 - b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client` if the value is not already populated.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the authorization window.
 - c. From the REST API UI, select **GET /assets**.
 - d. Select **Try it out**.
 - e. Select **Execute**.
 - f. From the response, copy the `"id"` from the `"storage"` section of the cluster you intend to check for upgrade readiness.



Do not use the `"parent"` value in this section because this is the management node's ID, not the storage cluster's ID.

```
"config": {},  
"credentialid": "12bbb2b2-f1be-123b-1234-12c3d4bc123e",  
"host_name": "SF_DEMO",  
"id": "12cc3a45-e6e7-8d91-a2bb-0bdb3456b789",  
"ip": "10.123.12.12",  
"parent": "d123ec42-456e-8912-ad3e-4bd56f4a789a",  
"sshcredentialid": null,  
"ssl_certificate": null
```

2. Run health checks on the storage cluster:

- Open the storage REST API UI on the management node:

```
https://<ManagementNodeIP>/storage/1/
```

- Select **Authorize** and complete the following:

- Enter the cluster user name and password.
- Enter the client ID as `mnnode-client` if the value is not already populated.
- Select **Authorize** to begin a session.
- Close the authorization window.

- Select **POST /health-checks**.

- Select **Try it out**.

- In the parameter field, enter the storage cluster ID obtained in Step 1.

```
{
  "config": {},
  "storageId": "123a45b6-1a2b-12a3-1234-1a2b34c567d8"
}
```

- Select **Execute** to run a health check on the specified storage cluster.

The response should indicate state as `initializing`:

```
{
  "_links": {
    "collection": "https://10.117.149.231/storage/1/health-checks",
    "log": "https://10.117.149.231/storage/1/health-checks/358f073f-896e-4751-ab7b-ccbb5f61f9fc/log",
    "self": "https://10.117.149.231/storage/1/health-checks/358f073f-896e-4751-ab7b-ccbb5f61f9fc"
  },
  "config": {},
  "dateCompleted": null,
  "dateCreated": "2020-02-21T22:11:15.476937+00:00",
  "healthCheckId": "358f073f-896e-4751-ab7b-ccbb5f61f9fc",
  "state": "initializing",
  "status": null,
  "storageId": "c6d124b2-396a-4417-8a47-df10d647f4ab",
  "taskId": "73f4df64-bda5-42c1-9074-b4e7843dbb77"
}
```

- g. Copy the healthCheckID that is part of response.
3. Verify the results of the health checks:
 - a. Select **GET /health-checks/{healthCheckId}**.
 - b. Select **Try it out**.

- c. Enter the health check ID in the parameter field.
- d. Select **Execute**.
- e. Scroll to the bottom of the response body.

If all health checks are successful, the return is similar to the following example:

```
"message": "All checks completed successfully.",  
"percent": 100,  
"timestamp": "2020-03-06T00:03:16.321621Z"
```

4. If the message return indicates that there were problems regarding cluster health, do the following:
 - a. Select **GET /health-checks/{healthCheckId}/log**
 - b. Select **Try it out**.
 - c. Enter the health check ID in the parameter field.
 - d. Select **Execute**.
 - e. Review any specific errors and obtain their associated KB article links.
 - f. Go to the specific KB article listed for each issue or perform the specified remedy.
 - g. If a KB is specified, complete the process described in the relevant KB article.
 - h. After you have resolved cluster issues, run **GET /health-checks/{healthCheckId}/log** again.

Use HealthTools to run Element storage health checks prior to upgrading storage

You can verify that the storage cluster is ready to be upgraded by using the `sfupgradecheck` command. This command verifies information such as pending nodes, disk space, and cluster faults.

If your management node is at a dark site, the upgrade readiness check needs the `metadata.json` file you downloaded during [HealthTools upgrades](#) to run successfully.

About this task

This procedure describes how to address upgrade checks that yield one of the following results:

- Running the `sfupgradecheck` command runs successfully. Your cluster is upgrade ready.
- Checks within the `sfupgradecheck` tool fail with an error message. Your cluster is not upgrade ready and additional steps are required.
- Your upgrade check fails with an error message that HealthTools is out-of-date.
- Your upgrade check fails because your management node is on a dark site.

Steps

1. Run the `sfupgradecheck` command:

```
sfupgradecheck -u <cluster-user-name> MVIP
```



For passwords that contain special characters, add a backslash (\) before each special character. For example, mypass!@1 should be entered as mypass\!@1.

Sample input command with sample output in which no errors appear and you are ready to upgrade:

```
sfupgradecheck -u admin 10.117.78.244
```

```
check_pending_nodes:
Test Description: Verify no pending nodes in cluster
More information:
https://kb.netapp.com/support/s/article/ka11A00000081tOQAQ/pendingnodes
check_cluster_faults:
Test Description: Report any cluster faults
check_root_disk_space:
Test Description: Verify node root directory has at least 12 GBs of
available disk space
Passed node IDs: 1, 2, 3
More information:
https://kb.netapp.com/support/s/article/ka11A00000081tTQAQ/SolidFire-Disk-space-error
check_mnode_connectivity:
Test Description: Verify storage nodes can communicate with management
node
Passed node IDs: 1, 2, 3
More information:
https://kb.netapp.com/support/s/article/ka11A00000081tYQAQ/mNodeconnectivity
check_files:
Test Description: Verify options file exists
Passed node IDs: 1, 2, 3
check_cores:
Test Description: Verify no core or dump files exists
Passed node IDs: 1, 2, 3
check_upload_speed:
Test Description: Measure the upload speed between the storage node and
the
management node
Node ID: 1 Upload speed: 90063.90 KBs/sec
Node ID: 3 Upload speed: 106511.44 KBs/sec
Node ID: 2 Upload speed: 85038.75 KBs/sec
```

2. If there are errors, additional actions are required. See the following sub-sections for details.

Your cluster is not upgrade ready

If you see an error message related to one of the health checks, follow these steps:

1. Review the `sfupgradecheck` error message.

Sample response:

The following tests failed:

check_root_disk_space:

Test Description: Verify node root directory has at least 12 GBs of available disk space

Severity: ERROR

Failed node IDs: 2

Remedy: Remove unneeded files from root drive

More information:

<https://kb.netapp.com/support/s/article/ka11A00000081tTQAQ/SolidFire-Disk-space-error>

check_pending_nodes:

Test Description: Verify no pending nodes in cluster

More information:

<https://kb.netapp.com/support/s/article/ka11A00000081tOQAQ/pendingnodes>

check_cluster_faults:

Test Description: Report any cluster faults

check_root_disk_space:

Test Description: Verify node root directory has at least 12 GBs of available disk space

Passed node IDs: 1, 3

More information:

<https://kb.netapp.com/support/s/article/ka11A00000081tTQAQ/SolidFire-Disk-space-error>

check_mnode_connectivity:

Test Description: Verify storage nodes can communicate with management node

Passed node IDs: 1, 2, 3

More information:

<https://kb.netapp.com/support/s/article/ka11A00000081tYQAQ/mNodeconnectivity>

check_files:

Test Description: Verify options file exists

Passed node IDs: 1, 2, 3

check_cores:

Test Description: Verify no core or dump files exists

Passed node IDs: 1, 2, 3

check_upload_speed:

Test Description: Measure the upload speed between the storage node and the management node

Node ID: 1 Upload speed: 86518.82 KBs/sec

Node ID: 3 Upload speed: 84112.79 KBs/sec

Node ID: 2 Upload speed: 93498.94 KBs/sec

In this example, node 1 is low on disk space. You can find more information in the [knowledge base \(KB\)](#) article listed in the error message.

HealthTools is out of date

If you see an error message indicating that HealthTools is not the latest version, follow these instructions:

1. Review the error message and note that the upgrade check fails.

Sample response:

```
sfupgradecheck failed: HealthTools is out of date:  
installed version: 2018.02.01.200  
latest version: 2020.03.01.09.  
The latest version of the HealthTools can be downloaded from:  
https://mysupport.netapp.com/NOW/cgi-bin/software/  
Or rerun with the -n option
```

2. Follow the instructions described in the response.

Your management node is on a dark site

1. Review the message and note that the upgrade check fails:

Sample response:

```
sfupgradecheck failed: Unable to verify latest available version of  
healthtools.
```

2. Download a [JSON file](#) from the NetApp Support Site on a computer that is not the management node and rename it to metadata.json.

3. Run the following command:

```
sfupgradecheck -l --metadata=<path-to-metadata-json>
```

4. For details, see additional [HealthTools upgrades](#) information for dark sites.

5. Verify that the HealthTools suite is up-to-date by running the following command:

```
sfupgradecheck -u <cluster-user-name> -p <cluster-password> MVIP
```

Storage health checks made by the service

Storage health checks make the following checks per cluster.

Check Name	Node/Cluster	Description
check_async_results	Cluster	Verifies that the number of asynchronous results in the database is below a threshold number.
check_cluster_faults	Cluster	Verifies that there are no upgrade blocking cluster faults (as defined in Element source).
check_upload_speed	Node	Measures the upload speed between the storage node and the management node.
connection_speed_check	Node	Verifies that nodes have connectivity to the management node serving upgrade packages and estimates connection speed.
check_cores	Node	Checks for kernel crash dump and core files on the node. The check fails for any crashes in a recent time period (threshold 7 days).
check_root_disk_space	Node	Verifies the root file system has sufficient free space to perform an upgrade.
check_var_log_disk_space	Node	Verifies that /var/log free space meets some percentage free threshold. If it does not, the check will rotate and purge older logs in order to fall under threshold. The check fails if it is unsuccessful at creating sufficient free space.
check_pending_nodes	Cluster	Verifies that there are no pending nodes on the cluster.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Upgrade Element software

To upgrade NetApp Element software, you can use the NetApp Hybrid Cloud Control UI, REST API, or the HealthTools suite of tools. Certain operations are suppressed during an Element software upgrade, such as adding and removing nodes, adding and removing drives, and commands associated with initiators, volume access groups, and virtual networks, among others.

What you'll need

- **Admin privileges:** You have storage cluster administrator permissions to perform the upgrade.
- **Valid upgrade path:** You have checked upgrade path information for the Element version you are

upgrading to and verified that the upgrade path is valid.

[NetApp KB: Upgrade matrix for storage clusters running NetApp Element Software \(login required\)](#)

- **System time sync:** You have ensured that the system time on all nodes is synced and that NTP is correctly configured for the storage cluster and nodes. Each node must be configured with a DNS nameserver in the per-node web UI ([https://\[IP address\]:442](https://[IP address]:442)) with no unresolved cluster faults related to time skew.
- **System ports:** If you are using NetApp Hybrid Cloud Control for upgrades, you have ensured that the necessary ports are open. See [Network ports](#) for more information.
- **Management node:** For NetApp Hybrid Cloud Control UI and API, the management node in your environment is running version 11.3.
- **Management services:** You have updated your management services bundle to the latest version.



You must upgrade to the latest management services bundle before upgrading your Element software to version 12. If you are updating your Element software to version 12.2, you need management services 2.14.60 or later to proceed.

- **Cluster health:** You have verified that the cluster is ready to be upgraded. See [Run Element storage health checks prior to upgrading storage](#).
- **Updated BMC for H610S nodes:** You have upgraded the BMC version for your H610S nodes. See the [release notes and upgrade instructions](#).
- **End User License Agreement (EULA):** Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to upgrade Element software:
 1. Open the IP address of the management node in a web browser:

`https://<ManagementNodeIP>`
 2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
 3. Select **Upgrade** near the top right of the interface.
 4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Upgrade options

Choose one of the following Element software upgrade options:

- [Use NetApp Hybrid Cloud Control UI to upgrade Element storage](#)
- [Use NetApp Hybrid Cloud Control API to upgrade Element storage](#)
- [Upgrade Element software at connected sites using HealthTools](#)
- [Upgrade Element software at dark sites using HealthTools](#)



If you are upgrading an H610S series node to Element 12.2 and the node is running a version of Element earlier than 11.8, you will need to perform additional upgrade steps ([phase 2](#)) for each storage node. If you are running Element 11.8 or later, the additional upgrade steps (phase 2) are not required.

Use NetApp Hybrid Cloud Control UI to upgrade Element storage

Using the NetApp Hybrid Cloud Control UI, you can upgrade a storage cluster.



For potential issues while upgrading storage clusters using NetApp Hybrid Cloud Control and their workarounds, see the [KB article](#).



The upgrade process takes approximately 30 minutes per node for non-H610S platforms.

Steps

1. Open the IP address of the management node in a web browser:

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

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

The **Storage** tab lists the storage clusters that are part of your installation. If a cluster is inaccessible by NetApp Hybrid Cloud Control, it will not be displayed on the **Upgrades** page.

5. Choose from the following options and perform the set of steps that are applicable to your cluster:

Option	Steps
<p>All clusters running Element 11.8 and later</p>	<ol style="list-style-type: none"> 1. Select Browse to upload the upgrade package that you downloaded. 2. Wait for the upload to complete. A progress bar shows the status of the upload. <p> The file upload will be lost if you navigate away from the browser window.</p> <p>An on-screen message is displayed after the file is successfully uploaded and validated. Validation might take several minutes. If you navigate away from the browser window at this stage, the file upload is preserved.</p> <ol style="list-style-type: none"> 3. Select Begin Upgrade. <p> The Upgrade Status changes during the upgrade to reflect the status of the process. It also changes in response to actions you take, such as pausing the upgrade, or if the upgrade returns an error. See Upgrade status changes.</p> <p> While the upgrade is in progress, you can leave the page and come back to it later to continue monitoring the progress. The page does not dynamically update status and current version if the cluster row is collapsed. The cluster row must be expanded to update the table or you can refresh the page.</p> <p>You can download logs after the upgrade is complete.</p>

Option	Steps
<p>You are upgrading an H610S cluster running Element version earlier than 11.8.</p>	<ol style="list-style-type: none"> 1. Select the drop-down arrow next to the cluster you are upgrading, and select from the upgrade versions available. 2. Select Begin Upgrade. After the upgrade is complete, the UI prompts you to perform phase 2 of the process. 3. Complete the additional steps required (phase 2) in the KB article, and acknowledge in the UI that you have completed phase 2. <p>You can download logs after the upgrade is complete. For information about the various upgrade status changes, see Upgrade status changes.</p>

Upgrade status changes

Here are the different states that the **Upgrade Status** column in the UI shows before, during, and after the upgrade process:

Upgrade state	Description
Up to Date	The cluster was upgraded to the latest Element version available.
Versions Available	Newer versions of Element and/or storage firmware are available for upgrade.
In Progress	The upgrade is in progress. A progress bar shows the upgrade status. On-screen messages also show node-level faults and display the node ID of each node in the cluster as the upgrade progresses. You can monitor the status of each node using the Element UI or the NetApp Element plug-in for vCenter Server UI.
Upgrade Pausing	You can choose to pause the upgrade. Depending on the state of the upgrade process, the pause operation can succeed or fail. You will see a UI prompt asking you to confirm the pause operation. To ensure that the cluster is in a safe spot before pausing an upgrade, it can take up to two hours for the upgrade operation to be completely paused. To resume the upgrade, select Resume .
Paused	You paused the upgrade. Select Resume to resume the process.

Upgrade state	Description
Error	An error has occurred during the upgrade. You can download the error log and send it to NetApp Support. After you resolve the error, you can return to the page, and select Resume . When you resume the upgrade, the progress bar goes backwards for a few minutes while the system runs the health check and checks the current state of the upgrade.
Complete with Follow-up	Only for H610S nodes upgrading from Element version earlier than 11.8. After phase 1 of the upgrade process is complete, this state prompts you to perform phase 2 of the upgrade (see the KB article). After you complete phase 2 and acknowledge that you have completed it, the status changes to Up to Date .

Use NetApp Hybrid Cloud Control API to upgrade Element storage

You can use APIs to upgrade storage nodes in a cluster to the latest Element software version. You can use an automation tool of your choice to run the APIs. The API workflow documented here uses the REST API UI available on the management node as an example.

Steps

1. Download the storage upgrade package to a device that is accessible to the management node; go to the NetApp HCI software [downloads page](#) and download the latest storage node image.

2. Upload the storage upgrade package to the management node:

a. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/package-repository/1/
```

b. Select **Authorize** and complete the following:

i. Enter the cluster user name and password.

ii. Enter the client ID as `mnode-client`.

iii. Select **Authorize** to begin a session.

iv. Close the authorization window.

c. From the REST API UI, select **POST /packages**.

d. Select **Try it out**.

e. Select **Browse** and select the upgrade package.

f. Select **Execute** to initiate the upload.

g. From the response, copy and save the package ID ("`id`") for use in a later step.

3. Verify the status of the upload.

a. From the REST API UI, select **GET /packages/{id}/status**.

b. Select **Try it out**.

c. Enter the package ID you copied in the previous step in **id**.

- d. Select **Execute** to initiate the status request.

The response indicates state as SUCCESS when complete.

4. Locate the storage cluster ID:

- a. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```

- b. Select **Authorize** and complete the following:

- i. Enter the cluster user name and password.
- ii. Enter the client ID as mnode-client.
- iii. Select **Authorize** to begin a session.
- iv. Close the authorization window.

- c. From the REST API UI, select **GET /installations**.

- d. Select **Try it out**.

- e. Select **Execute**.

- f. From the response, copy the installation asset ID ("id").

- g. From the REST API UI, select **GET /installations/{id}**.

- h. Select **Try it out**.

- i. Paste the installation asset ID into the **id** field.

- j. Select **Execute**.

- k. From the response, copy and save the storage cluster ID ("id") of the cluster you intend to upgrade for use in a later step.

5. Run the storage upgrade:

- a. Open the storage REST API UI on the management node:

```
https://<ManagementNodeIP>/storage/1/
```

- b. Select **Authorize** and complete the following:

- i. Enter the cluster user name and password.
- ii. Enter the client ID as mnode-client.
- iii. Select **Authorize** to begin a session.
- iv. Close the authorization window.

- c. Select **POST /upgrades**.

- d. Select **Try it out**.

- e. Enter the upgrade package ID in the parameter field.

- f. Enter the storage cluster ID in the parameter field.

The payload should look similar to the following example:

```
{  
  "config": {},  
  "packageId": "884f14a4-5a2a-11e9-9088-6c0b84e211c4",  
  "storageId": "884f14a4-5a2a-11e9-9088-6c0b84e211c4"  
}
```

g. Select **Execute** to initiate the upgrade.

The response should indicate the state as `initializing`:

```
{  
  "_links": {  
    "collection": "https://localhost:442/storage/upgrades",  
    "self": "https://localhost:442/storage/upgrades/3fa85f64-1111-  
4562-b3fc-2c963f66abc1",  
    "log": "https://localhost:442/storage/upgrades/3fa85f64-1111-4562-  
b3fc-2c963f66abc1/log"  
  },  
  "storageId": "114f14a4-1a1a-11e9-9088-6c0b84e200b4",  
  "upgradeId": "334f14a4-1a1a-11e9-1055`-6c0b84e2001b4",  
  "packageId": "774f14a4-1a1a-11e9-8888-6c0b84e200b4",  
  "config": {},  
  "state": "initializing",  
  "status": {  
    "availableActions": [  
      "string"  
    ],  
    "message": "string",  
    "nodeDetails": [  
      {  
        "message": "string",  
        "step": "NodePreStart",  
        "nodeID": 0,  
        "numAttempt": 0  
      }  
    ],  
    "percent": 0,  
    "step": "ClusterPreStart",  
    "timestamp": "2020-04-21T22:10:57.057Z",  
    "failedHealthChecks": [  
      {  
        "checkID": 0,  
        "name": "string",  
        "status": "string"  
      }  
    ]  
  }  
}
```

```

        "displayName": "string",
        "passed": true,
        "kb": "string",
        "description": "string",
        "remedy": "string",
        "severity": "string",
        "data": {},
        "nodeID": 0
    }
]
},
"taskId": "123f14a4-1a1a-11e9-7777-6c0b84e123b2",
"dateCompleted": "2020-04-21T22:10:57.057Z",
"dateCreated": "2020-04-21T22:10:57.057Z"
}

```

- h. Copy the upgrade ID ("upgradeId") that is part of the response.
6. Verify the upgrade progress and results:
- a. Select **GET /upgrades/{upgradeld}**.
 - b. Select **Try it out**.
 - c. Enter the upgrade ID from the previous step in **upgradeld**.
 - d. Select **Execute**.
 - e. Do one of the following if there are problems or special requirements during the upgrade:

Option	Steps
<p>You need to correct cluster health issues due to <code>failedHealthChecks</code> message in the response body.</p>	<ol style="list-style-type: none"> 1. Go to the specific KB article listed for each issue or perform the specified remedy. 2. If a KB is specified, complete the process described in the relevant KB article. 3. After you have resolved cluster issues, reauthenticate if needed and select PUT /upgrades/{upgradeld}. 4. Select Try it out. 5. Enter the upgrade ID from the previous step in upgradeld. 6. Enter "action": "resume" in the request body. <div data-bbox="913 686 1493 865" style="border: 1px solid #ccc; padding: 10px; width: fit-content; margin: 10px auto;"> <pre>{ "action": "resume" }</pre> </div> <ol style="list-style-type: none"> 7. Select Execute.
<p>You need to pause the upgrade because the maintenance window is closing or for another reason.</p>	<ol style="list-style-type: none"> 1. Reauthenticate if needed and select PUT /upgrades/{upgradeld}. 2. Select Try it out. 3. Enter the upgrade ID from the previous step in upgradeld. 4. Enter "action": "pause" in the request body. <div data-bbox="913 1298 1493 1478" style="border: 1px solid #ccc; padding: 10px; width: fit-content; margin: 10px auto;"> <pre>{ "action": "pause" }</pre> </div> <ol style="list-style-type: none"> 5. Select Execute.

Option	Steps
<p>If you are upgrading an H610S cluster running an Element version earlier than 11.8, you see the state <code>finishedNeedsAck</code> in the response body. You need to perform additional upgrade steps (phase 2) for each H610S storage node.</p>	<ol style="list-style-type: none"> 1. See [Upgrading H610S storage nodes to Element 12.2 or later (phase 2)] and complete the process for each node. 2. Reauthenticate if needed and select PUT /upgrades/{upgradeld}. 3. Select Try it out. 4. Enter the upgrade ID from the previous step in <code>upgradeld</code>. 5. Enter <code>"action": "acknowledge"</code> in the request body. <div style="border: 1px solid #ccc; padding: 10px; width: fit-content; margin-left: 20px;"> <pre>{ "action": "acknowledge" }</pre> </div> <ol style="list-style-type: none"> 6. Select Execute.

- f. Run the **GET /upgrades/{upgradeld}** API multiple times, as needed, until the process is complete.

During the upgrade, the `status` indicates `running` if no errors are encountered. As each node is upgraded, the `step` value changes to `NodeFinished`.

The upgrade has finished successfully when the `percent` value is 100 and the `state` indicates `finished`.

What happens if an upgrade fails using NetApp Hybrid Cloud Control

If a drive or node fails during an upgrade, the Element UI will show cluster faults. The upgrade process does not proceed to the next node, and waits for the cluster faults to resolve. The progress bar in the UI shows that the upgrade is waiting for the cluster faults to resolve. At this stage, selecting **Pause** in the UI will not work, because the upgrade waits for the cluster to be healthy. You will need to engage NetApp Support to assist with the failure investigation.

NetApp Hybrid Cloud Control has a pre-set three-hour waiting period, during which one of the following scenarios can happen:

- The cluster faults get resolved within the three-hour window, and upgrade resumes. You do not need to take any action in this scenario.
- The problem persists after three hours, and the upgrade status shows **Error** with a red banner. You can resume the upgrade by selecting **Resume** after the problem is resolved.
- NetApp Support has determined that the upgrade needs to be temporarily aborted to take corrective action before the three-hour window. Support will use the API to abort the upgrade.

 Aborting the cluster upgrade while a node is being updated might result in the drives being ungracefully removed from the node. If the drives are ungracefully removed, adding the drives back during an upgrade will require manual intervention by NetApp Support. The node might be taking longer to do firmware updates or post update syncing activities. If the upgrade progress seems stalled, contact NetApp Support for assistance.

Upgrade Element software at connected sites using HealthTools

Steps

1. Download the storage upgrade package; go to the NetApp HCI software [downloads page](#) and download the latest storage node image to a device that is not the management node.



You need the latest version of HealthTools to upgrade Element storage software.

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

When you upload the ISO file, make sure that the name of the file does not change, otherwise later steps will fail.

3. **Optional:** Download the ISO from the management node to the cluster nodes before the upgrade.

This step reduces the upgrade time by pre-staging the ISO on the storage nodes and running additional internal checks to ensure that the cluster is in a good state to be upgraded. Performing this operation will not put the cluster into "upgrade" mode or restrict any of the cluster operations.

```
sfinstall <MVIP> -u <cluster_username> <path-toinstall-file-ISO> --stage
```



Omit the password from the command line to allow sfinstall to prompt for the information. For passwords that contain special characters, add a backslash (\) before each special character. For example, mypass!@1 should be entered as mypass\\!\\@.

Example

See the following sample input:

```
sfinstall 10.117.0.244 -u admin /tmp/solidfire-rtfisodium-11.0.0.345.iso  
--stage
```

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

```
sfinstall 10.117.0.244 -u admin
/ttmp/solidfire-rtfisodium-11.0.0.345.iso 2018-10-01 16:52:15:
Newer version of sfinstall available.
This version: 2018.09.01.130, latest version: 2018.06.05.901.
The latest version of the HealthTools can be downloaded from:
https://mysupport.netapp.com/NOW/cgi-bin/software/
or rerun with --skip-version-check
```

See the following sample excerpt from a successful pre-stage operation:



When staging completes, the message will display **Storage Node Upgrade Staging Successful** after the upgrade event.

```
flabv0004 ~ # sfinstall -u admin
10.117.0.87 solidfire-rtfi-sodium-patch3-11.3.0.14171.iso --stage
2019-04-03 13:19:58: sfinstall Release Version: 2019.01.01.49 Management
Node Platform:
Ember Revision: 26b042c3e15a Build date: 2019-03-12 18:45
2019-04-03 13:19:58: Checking connectivity to MVIP 10.117.0.87
2019-04-03 13:19:58: Checking connectivity to node 10.117.0.86
2019-04-03 13:19:58: Checking connectivity to node 10.117.0.87
...
2019-04-03 13:19:58: Successfully connected to cluster and all nodes
...
2019-04-03 13:20:00: Do you want to continue? ['Yes', 'No']: Yes
...
2019-04-03 13:20:55: Staging install pack on cluster nodes
2019-04-03 13:20:55: newVersion: 11.3.0.14171
2019-04-03 13:21:01: nodeToStage: nlabp2814, nlabp2815, nlabp2816,
nlabp2813
2019-04-03 13:21:02: Staging Node nlabp2815 mip=[10.117.0.87] nodeID=[2]
(1 of 4 nodes)
2019-04-03 13:21:02: Node Upgrade serving image at
http://10.117.0.204/rtfi/solidfire-rtfisodium-
patch3-11.3.0.14171/filesystem.squashfs
...
2019-04-03 13:25:40: Staging finished. Repeat the upgrade command
without the --stage option to start the upgrade.
```

The staged ISOs will be automatically deleted after the upgrade completes. However, if the upgrade has not started and needs to be rescheduled, ISOs can be manually de-staged using the command:

```
sfinstall <MVIP> -u <cluster_username> --destage
```

After the upgrade has started, the de-stage option is no longer available.

4. Start the upgrade with the `sfinstall` command and the path to the ISO file:

```
sfinstall <MVIP> -u <cluster_username> <path-toinstall-file-ISO>
```

Example

See the following sample input command:

```
sfinstall 10.117.0.244 -u admin /tmp/solidfire-rtfi-sodium-11.0.0.345.iso
```

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

```
sfinstall 10.117.0.244 -u admin /tmp/solidfire-rtfi-sodium-11.0.0.345.iso
2018-10-01 16:52:15: Newer version of sfinstall available.
This version: 2018.09.01.130, latest version: 2018.06.05.901.
The latest version of the HealthTools can be downloaded from:
https://mysupport.netapp.com/NOW/cgi-bin/software/ or rerun with --skip
--version-check
```

See the following sample excerpt from a successful upgrade. Upgrade events can be used to monitor the progress of the upgrade.

```
# sfinstall 10.117.0.161 -u admin solidfire-rtfi-sodium-11.0.0.761.iso
2018-10-11 18:28
Checking connectivity to MVIP 10.117.0.161
Checking connectivity to node 10.117.0.23
Checking connectivity to node 10.117.0.24
...
Successfully connected to cluster and all nodes
#####
You are about to start a new upgrade
10.117.0.161
10.3.0.161
solidfire-rtfi-sodium-11.0.0.761.iso
Nodes:
10.117.0.23 nlabp1023 SF3010 10.3.0.161
10.117.0.24 nlabp1025 SF3010 10.3.0.161
10.117.0.26 nlabp1027 SF3010 10.3.0.161
10.117.0.28 nlabp1028 SF3010 10.3.0.161
#####
Do you want to continue? ['Yes', 'No']: yes
...
```

```
Watching for new network faults. Existing fault IDs are set([]).
Checking for legacy network interface names that need renaming
Upgrading from 10.3.0.161 to 11.0.0.761 upgrade method=rtfi
Waiting 300 seconds for cluster faults to clear
Waiting for caches to fall below threshold
...
Installing mip=[10.117.0.23] nodeID=[1] (1 of 4 nodes)
Starting to move primaries.
Loading volume list
Moving primary slice=[7] away from mip[10.117.0.23] nodeID[1] ssid[11]
to new ssid[15]
Moving primary slice=[12] away from mip[10.117.0.23] nodeID[1] ssid[11]
to new ssid[15]
...
Installing mip=[10.117.114.24] nodeID=[2] (2 of 4 nodes)
Starting to move primaries.
Loading volume list
Moving primary slice=[5] away from mip[10.117.114.24] nodeID[2] ssid[7]
to new ssid[11]
...
Install of solidfire-rtfi-sodium-11.0.0.761 complete.
Removing old software
No staged builds present on nodeID=[1]
No staged builds present on nodeID=[2]
...
Starting light cluster block service check
```

 If you are upgrading an H610S series node to Element 12.2 and the node is running a version of Element earlier than 11.8, you will need to perform additional upgrade steps ([phase 2](#)) for each storage node. If you are running Element 11.8 or later, the additional upgrade steps (phase 2) are not required.

Upgrade Element software at dark sites using HealthTools

You can use the HealthTools suite of tools to update NetApp Element software at a dark site.

What you'll need

1. Go to the NetApp HCI software [downloads page](#).
 2. Select the correct software release and download the latest storage node image to a computer that is not the management node.
-  You need the latest version of HealthTools to upgrade Element storage software.
3. Download this [JSON file](#) from the NetApp Support Site on a computer that is not the management node and rename it to `metadata.json`.
 4. Copy the ISO file to the management node in an accessible location like `/tmp`.



You can do this by using, for example, SCP. When you upload the ISO file, make sure that the name of the file does not change, otherwise later steps will fail.

Steps

1. Run the `sfupdate-healthtools` command:

```
sfupdate-healthtools <path-to-healthtools-package>
```

2. Check the installed version:

```
sfupdate-healthtools -v
```

3. Check the latest version against the metadata JSON file:

```
sfupdate-healthtools -l --metadata=<path-to-metadata-json>
```

4. Ensure that the cluster is ready:

```
sudo sfupgradecheck -u <cluster_username> -p <cluster_password> MVIP  
--metadata=<path-to-metadata-json>
```

5. Run the `sfinstall` command with the path to the ISO file and the metadata JSON file:

```
sfinstall -u <cluster_username> <MVIP> <path-toinstall-file-ISO>  
--metadata=<path-to-metadata-json-file>
```

See the following sample input command:

```
sfinstall -u admin 10.117.78.244 /tmp/solidfire-rtfi-11.3.0.345.iso  
--metadata=/tmp/metadata.json
```

Optional You can add the `--stage` flag to the `sfinstall` command to pre-stage the upgrade in advance.



If you are upgrading an H610S series node to Element 12.2 and the node is running a version of Element earlier than 11.8, you will need to perform additional upgrade steps ([phase 2](#)) for each storage node. If you are running Element 11.8 or later, the additional upgrade steps ([phase 2](#)) are not required.

What happens if an upgrade fails using HealthTools

If the software upgrade fails, you can pause the upgrade.



You should pause an upgrade only with Ctrl-C. This enables the system to clean itself up.

When `sfinstall` waits for cluster faults to clear and if any failure causes the faults to remain, `sfinstall` will not proceed to the next node.

Steps

1. You should stop `sfinstall` with Ctrl+C.
2. Contact NetApp Support to assist with the failure investigation.
3. Resume the upgrade with the same `sfinstall` command.
4. When an upgrade is paused by using Ctrl+C, if the upgrade is currently upgrading a node, choose one of these options:
 - **Wait:** Allow the currently upgrading node to finish before resetting the cluster constants.
 - **Continue:** Continue the upgrade, which cancels the pause.
 - **Abort:** Reset the cluster constants and abort the upgrade immediately.



Aborting the cluster upgrade while a node is being updated might result in the drives being ungracefully removed from the node. If the drives are ungracefully removed, adding the drives back during an upgrade will require manual intervention by NetApp Support. The node might be taking longer to do firmware updates or post update syncing activities. If the upgrade progress seems stalled, contact NetApp Support for assistance.

Upgrading H610S storage nodes to Element 12.2 (phase 2)

If you are upgrading an H610S series node to Element 12.2 and the node is running a version of Element earlier than 11.8, the upgrade process involves two phases.

Phase 1, which is performed first, follows the same steps as the standard upgrade to Element 12.2 process. It installs Element Software and all 5 firmware updates in a rolling fashion across the cluster one node at a time. Due to the firmware payload, the process is estimated to take approximately 1.5 to 2 hours per H610S node, including a single cold-boot cycle at the end of the upgrade for each node.

Phase 2 involves completing steps to perform a complete node shutdown and power disconnect for each H610S node that are described in a required [KB](#). This phase is estimated to take approximately one hour per H610S node.



After you complete phase 1, four of the five firmware updates are activated during the cold boot on each H610S node; however, the Complex Programmable Logic Device (CPLD) firmware requires a complete power disconnect and reconnect to fully install. The CPLD firmware update protects against NVDIMM failures and metadata drive eviction during future reboots or power cycles. This power reset is estimated to take approximately one hour per H610S node. It requires shutting down the node, removing power cables or disconnecting power via a smart PDU, waiting approximately 3 minutes, and reconnecting power.

What you'll need

- You have completed phase 1 of the H610S upgrade process and have upgraded your storage nodes using

one the standard Element storage upgrade procedures.



Phase 2 requires on-site personnel.

Steps

1. (Phase 2) Complete the power reset process required for each H610S node in the cluster:



If the cluster also has non-H610S nodes, these non-H610S nodes are exempt from phase 2 and do not need to be shut down or have their power disconnected.

- a. Contact NetApp Support for assistance and to schedule this upgrade.
- b. Follow the phase 2 upgrade procedure in this [KB](#) that is required to complete an upgrade for each H610S node.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Upgrade storage firmware

Starting with Element 12.0 and management services version 2.14, you can perform firmware-only upgrades on your H-series storage nodes using the NetApp Hybrid Cloud Control UI and REST API. This procedure does not upgrade Element software and enables you to upgrade storage firmware outside of a major Element release.

What you'll need

- **Admin privileges:** You have storage cluster administrator permissions to perform the upgrade.
- **System time sync:** You have ensured that the system time on all nodes is synced and that NTP is correctly configured for the storage cluster and nodes. Each node must be configured with a DNS nameserver in the per-node web UI ([https://\[IP address\]:442](https://[IP address]:442)) with no unresolved cluster faults related to time skew.
- **System ports:** If you are using NetApp Hybrid Cloud Control for upgrades, you have ensured that the necessary ports are open. See [Network ports](#) for more information.
- **Management node:** For NetApp Hybrid Cloud Control UI and API, the management node in your environment is running version 11.3.
- **Management services:** You have updated your management services bundle to the latest version.



For H610S storage nodes running Element software version 12.0, you should apply D-patch SUST-909 before you upgrade to storage firmware bundle 2.27. Contact NetApp Support to obtain the D-patch before you upgrade. See [Storage Firmware Bundle 2.27 Release Notes](#).



You must upgrade to the latest management services bundle before upgrading the firmware on your storage nodes. If you are updating your Element software to version 12.2, you need management services 2.14.60 or later to proceed.

- **Cluster health:** You have run health checks. See [Run Element storage health checks prior to upgrading storage](#).
- **Updated BMC for H610S nodes:** You have upgraded the BMC version for your H610S nodes. See

release notes and upgrade instructions.



For a complete matrix of firmware and driver firmware for your hardware, see [supported firmware versions for NetApp HCI storage nodes](#).

- **End User License Agreement (EULA):** Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to upgrade storage firmware:

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.

3. Select **Upgrade** near the top right of the interface.

4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Upgrade options

Choose one of the following storage firmware upgrade options:

- [Use NetApp Hybrid Cloud Control UI to upgrade storage firmware](#)
- [Use NetApp Hybrid Cloud Control API to upgrade storage firmware](#)

Use NetApp Hybrid Cloud Control UI to upgrade storage firmware

You can use the NetApp Hybrid Cloud Control UI to upgrade the firmware of the storage nodes in your cluster.

What you'll need

- If your management node is not connected to the internet, you have [downloaded the storage firmware bundle](#).



For potential issues while upgrading storage clusters using NetApp Hybrid Cloud Control and their workarounds, see the [KB article](#).



The upgrade process takes approximately 30 minutes per node.

Steps

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.

3. Select **Upgrade** near the top right of the interface.

4. On the **Upgrades** page, select **Storage**.

The **Storage** tab lists the storage clusters that are part of your installation. If a cluster is inaccessible by NetApp Hybrid Cloud Control, it will not be displayed on the **Upgrades** page. If you have clusters running Element 12.0 or later, you will see the current firmware bundle version listed for these clusters. If the nodes in a single cluster have different firmware versions on them or as the upgrade progresses, you will see **Multiple** in the **Current Firmware Bundle Version** column. You can select **Multiple** to navigate to the **Nodes** page to compare firmware versions. If all your clusters are running Element versions earlier than 12.0, you will not see any information about firmware bundle version numbers. This information is also available on the **Nodes** page. See [View your inventory](#).

If the cluster is up to date and/or no upgrade packages are available, the **Element** and **Firmware Only** tabs are not displayed. These tabs are also not displayed when an upgrade is in progress. If the **Element** tab is displayed, but not the **Firmware Only** tab, no firmware packages are available.

5. Select the drop-down arrow next to the cluster you are upgrading.
6. Select **Browse** to upload the upgrade package that you downloaded.
7. Wait for the upload to complete. A progress bar shows the status of the upload.



The file upload will be lost if you navigate away from the browser window.

An on-screen message is displayed after the file is successfully uploaded and validated. Validation might take several minutes. If you navigate away from the browser window at this stage, the file upload is preserved.

8. Select **Firmware Only**, and select from the upgrade versions available.
9. Select **Begin Upgrade**.



The **Upgrade Status** changes during the upgrade to reflect the status of the process. It also changes in response to actions you take, such as pausing the upgrade, or if the upgrade returns an error. See [Upgrade status changes](#).



While the upgrade is in progress, you can leave the page and come back to it later to continue monitoring the progress. The page does not dynamically update status and current version if the cluster row is collapsed. The cluster row must be expanded to update the table or you can refresh the page.

You can download logs after the upgrade is complete.

Upgrade status changes

Here are the different states that the **Upgrade Status** column in the UI shows before, during, and after the upgrade process:

Upgrade state	Description
Up to Date	The cluster was upgraded to the latest Element version available or the firmware was upgraded to the latest version.

Upgrade state	Description
Unable to Detect	This status is displayed when the storage service API returns an upgrade status that is not in the enumerated list of possible upgrade statuses.
Versions Available	Newer versions of Element and/or storage firmware are available for upgrade.
In Progress	The upgrade is in progress. A progress bar shows the upgrade status. On-screen messages also show node-level faults and display the node ID of each node in the cluster as the upgrade progresses. You can monitor the status of each node using the Element UI or the NetApp Element plug-in for vCenter Server UI.
Upgrade Pausing	You can choose to pause the upgrade. Depending on the state of the upgrade process, the pause operation can succeed or fail. You will see a UI prompt asking you to confirm the pause operation. To ensure that the cluster is in a safe spot before pausing an upgrade, it can take up to two hours for the upgrade operation to be completely paused. To resume the upgrade, select Resume .
Paused	You paused the upgrade. Select Resume to resume the process.
Error	An error has occurred during the upgrade. You can download the error log and send it to NetApp Support. After you resolve the error, you can return to the page, and select Resume . When you resume the upgrade, the progress bar goes backwards for a few minutes while the system runs the health check and checks the current state of the upgrade.

What happens if an upgrade fails using NetApp Hybrid Cloud Control

If a drive or node fails during an upgrade, the Element UI will show cluster faults. The upgrade process does not proceed to the next node, and waits for the cluster faults to resolve. The progress bar in the UI shows that the upgrade is waiting for the cluster faults to resolve. At this stage, selecting **Pause** in the UI will not work, because the upgrade waits for the cluster to be healthy. You will need to engage NetApp Support to assist with the failure investigation.

NetApp Hybrid Cloud Control has a pre-set three-hour waiting period, during which one of the following scenarios can happen:

- The cluster faults get resolved within the three-hour window, and upgrade resumes. You do not need to take any action in this scenario.
- The problem persists after three hours, and the upgrade status shows **Error** with a red banner. You can resume the upgrade by selecting **Resume** after the problem is resolved.
- NetApp Support has determined that the upgrade needs to be temporarily aborted to take corrective action before the three-hour window. Support will use the API to abort the upgrade.

 Aborting the cluster upgrade while a node is being updated might result in the drives being ungracefully removed from the node. If the drives are ungracefully removed, adding the drives back during an upgrade will require manual intervention by NetApp Support. The node might be taking longer to do firmware updates or post update syncing activities. If the upgrade progress seems stalled, contact NetApp Support for assistance.

Use NetApp Hybrid Cloud Control API to upgrade storage firmware

You can use APIs to upgrade storage nodes in a cluster to the latest Element software version. You can use an automation tool of your choice to run the APIs. The API workflow documented here uses the REST API UI available on the management node as an example.

Steps

1. Download the latest storage firmware upgrade package to a device that is accessible to the management node; go to the [Element software storage firmware bundle page](#) and download the latest storage firmware image.
2. Upload the storage firmware upgrade package to the management node:
 - a. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/package-repository/1/
```
 - b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client`.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the authorization window.
 - c. From the REST API UI, select **POST /packages**.
 - d. Select **Try it out**.
 - e. Select **Browse** and select the upgrade package.
 - f. Select **Execute** to initiate the upload.
 - g. From the response, copy and save the package ID ("`id`") for use in a later step.
3. Verify the status of the upload.
 - a. From the REST API UI, select **GET /packages/{id}/status**.
 - b. Select **Try it out**.
 - c. Enter the firmware package ID you copied in the previous step in **id**.
 - d. Select **Execute** to initiate the status request.
4. Locate the installation asset ID:
 - a. Open the management node REST API UI on the management node:

The response indicates `state` as `SUCCESS` when complete.

4. Locate the installation asset ID:
 - a. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```

- b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as mnode-client.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the authorization window.
- c. From the REST API UI, select **GET /installations**.
- d. Select **Try it out**.
- e. Select **Execute**.
- f. From the response, copy the installation asset ID (**id**).

```
"id": "abcd01e2-xx00-4ccf-11ee-11f111xx9a0b",
"management": {
  "errors": [],
  "inventory": {
    "authoritativeClusterMvip": "10.111.111.111",
    "bundleVersion": "2.14.19",
    "managementIp": "10.111.111.111",
    "version": "1.4.12"
  }
}
```

- g. From the REST API UI, select **GET /installations/{id}**.
- h. Select **Try it out**.
- i. Paste the installation asset ID into the **id** field.
- j. Select **Execute**.
- k. From the response, copy and save the storage cluster ID ("id") of the cluster you intend to upgrade for use in a later step.

```
"storage": {
  "errors": [],
  "inventory": {
    "clusters": [
      {
        "clusterUuid": "a1bd1111-4f1e-46zz-ab6f-0a1111b1111x",
        "id": "a1bd1111-4f1e-46zz-ab6f-a1a1a111b012",
        ...
      }
    ]
  }
}
```

5. Run the storage firmware upgrade:
 - a. Open the storage REST API UI on the management node:

```
https://<ManagementNodeIP>/storage/1/
```

b. Select **Authorize** and complete the following:

- i. Enter the cluster user name and password.
- ii. Enter the client ID as mnode-client.
- iii. Select **Authorize** to begin a session.
- iv. Close the window.

c. Select **POST /upgrades**.

d. Select **Try it out**.

- e. Enter the upgrade package ID in the parameter field.
- f. Enter the storage cluster ID in the parameter field.
- g. Select **Execute** to initiate the upgrade.

The response should indicate state as **initializing**:

```
{
  "_links": {
    "collection": "https://localhost:442/storage/upgrades",
    "self": "https://localhost:442/storage/upgrades/3fa85f64-1111-4562-b3fc-2c963f66abc1",
    "log": "https://localhost:442/storage/upgrades/3fa85f64-1111-4562-b3fc-2c963f66abc1/log"
  },
  "storageId": "114f14a4-1a1a-11e9-9088-6c0b84e200b4",
  "upgradeId": "334f14a4-1a1a-11e9-1055-6c0b84e2001b4",
  "packageId": "774f14a4-1a1a-11e9-8888-6c0b84e200b4",
  "config": {},
  "state": "initializing",
  "status": {
    "availableActions": [
      "string"
    ],
    "message": "string",
    "nodeDetails": [
      {
        "message": "string",
        "step": "NodePreStart",
        "nodeID": 0,
        "numAttempt": 0
      }
    ],
    "percent": 0
  }
}
```

```

"step": "ClusterPreStart",
"timestamp": "2020-04-21T22:10:57.057Z",
"failedHealthChecks": [
  {
    "checkID": 0,
    "name": "string",
    "displayName": "string",
    "passed": true,
    "kb": "string",
    "description": "string",
    "remedy": "string",
    "severity": "string",
    "data": {},
    "nodeID": 0
  }
],
},
"taskId": "123f14a4-1a1a-11e9-7777-6c0b84e123b2",
"dateCompleted": "2020-04-21T22:10:57.057Z",
"dateCreated": "2020-04-21T22:10:57.057Z"
}

```

- h. Copy the upgrade ID ("upgradeId") that is part of the response.
6. Verify the upgrade progress and results:
 - a. Select **GET /upgrades/{upgradeld}**.
 - b. Select **Try it out**.
 - c. Enter the upgrade ID from the previous step in **upgradeld**.
 - d. Select **Execute**.
 - e. Do one of the following if there are problems or special requirements during the upgrade:

Option	Steps
<p>You need to correct cluster health issues due to <code>failedHealthChecks</code> message in the response body.</p>	<ol style="list-style-type: none"> 1. Go to the specific KB article listed for each issue or perform the specified remedy. 2. If a KB is specified, complete the process described in the relevant KB article. 3. After you have resolved cluster issues, reauthenticate if needed and select PUT /upgrades/{upgradeld}. 4. Select Try it out. 5. Enter the upgrade ID from the previous step in <code>upgradeld</code>. 6. Enter <code>"action": "resume"</code> in the request body. <div data-bbox="915 692 1491 868" style="border: 1px solid #ccc; padding: 10px; width: fit-content; margin: 10px auto;"> <pre>{ "action": "resume" }</pre> </div> <ol style="list-style-type: none"> 7. Select Execute.
<p>You need to pause the upgrade because the maintenance window is closing or for another reason.</p>	<ol style="list-style-type: none"> 1. Reauthenticate if needed and select PUT /upgrades/{upgradeld}. 2. Select Try it out. 3. Enter the upgrade ID from the previous step in <code>upgradeld</code>. 4. Enter <code>"action": "pause"</code> in the request body. <div data-bbox="915 1305 1491 1480" style="border: 1px solid #ccc; padding: 10px; width: fit-content; margin: 10px auto;"> <pre>{ "action": "pause" }</pre> </div> <ol style="list-style-type: none"> 5. Select Execute.

f. Run the **GET /upgrades/{upgradeld}** API multiple times, as needed, until the process is complete.

During the upgrade, the `status` indicates `running` if no errors are encountered. As each node is upgraded, the `step` value changes to `NodeFinished`.

The upgrade has finished successfully when the `percent` value is 100 and the `state` indicates `finished`.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Upgrade a management node

You can upgrade your management node to management node version 12.2 from version 11.0 or later.



Upgrading the management node operating system is no longer required to upgrade Element software on the storage cluster. If the management node is version 11.3 or higher, you can simply upgrade the management services to the latest version to perform Element upgrades using NetApp Hybrid Cloud Control. Follow the management node upgrade procedure for your scenario if you would like to upgrade the management node operating system for other reasons, such as security remediation.

What you'll need

- The vCenter Plug-in 4.4 or later requires a management node 11.3 or later that is created with modular architecture and provides individual services.

Upgrade options

Choose one of the following management node upgrade options:

- If you are upgrading from management node 12.0:
[Upgrade a management node to version 12.2 from 12.0](#)
- If you are upgrading from management node 11.3, 11.5, 11.7, or 11.8:
[Upgrade a management node to version 12.2 from 11.3 through 11.8](#)
- If you are upgrading from management node 11.0 or 11.1:
[Upgrade a management node to version 12.2 from 11.1 or 11.0](#)
- If you are upgrading from a management node version 10.x:
[Migrating from management node version 10.x to 11.x](#)

Choose this option if you have **sequentially** updated (1) your management services version and (2) your Element storage version and you want to **keep** your existing management node:



If you do not sequentially update your management services followed by Element storage, you cannot reconfigure reauthentication using this procedure. Follow the appropriate upgrade procedure instead.

- If you are keeping existing management node:
[Reconfigure authentication using the management node REST API](#)

Upgrade a management node to version 12.2 from 12.0

You can perform an in-place upgrade of the management node from version 12.0 to version 12.2 without needing to provision a new management node virtual machine.



The Element 12.2 management node is an optional upgrade. It is not required for existing deployments.

What you'll need

- The management node you are intending to upgrade is version 12.0 and uses IPv4 networking. The management node version 12.2 does not support IPv6.



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 using NetApp Hybrid Cloud Control (HCC). You can access HCC from the following IP: <code><a href="https://<ManagementNodeIP>" class="bare">https://<ManagementNodeIP></code>
- If you are updating your management node to version 12.2, you need management services 2.14.60 or later to proceed.
- You have configured an additional network adapter (if required) using the instructions for [configuring an additional storage NIC](#).



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.

- Storage nodes are running Element 11.3 or later.

Steps

1. Configure the management node VM RAM:
 - a. Power off the management node VM.
 - b. Change the RAM of the management node VM from 12GB to 24GB RAM.
 - c. Power on the management node VM.
2. Log in to the management node virtual machine using SSH or console access.
3. Download the [management node ISO](#) for NetApp HCI from the NetApp Support Site to the management node virtual machine.



The name of the ISO is similar to solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso

4. Check the integrity of the download by running md5sum on the downloaded file and compare 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-<Element release>-patchX-XX.X.X.XXXX.iso
```

5. 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-<Element release>-patchX-XX.X.X.XXXX.iso> /mnt
```

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

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

```
sudo umount /mnt
```

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

```
sudo rm <path to iso>/solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso
```

8. (For configurations without persistent volumes only) Copy the contents of the container folder for backup:

```
sudo cp -r /var/lib/docker/volumes /sf/etc/mnode
```

9. On the management node that you are upgrading, 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/rtdi/bin/sfrtdi_inplace  
file:///upgrade/casper/filesystem.squashfs sf_upgrade=1
```

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

10. (For configurations without persistent volumes only) Move the contents of the container folder back to original location:

```
sudo su  
mv /sf/etc/mnode/volumes/* /var/lib/docker/volumes/
```

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



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>
```



If you had previously disabled SSH functionality on the management node, you need to [disable SSH again](#) on the recovered management node. SSH capability that provides [NetApp Support remote support tunnel \(RST\) session access](#) is enabled on the management node by default.

Upgrade a management node to version 12.2 from 11.3 through 11.8

You can perform an in-place upgrade of the management node from version 11.3, 11.5, 11.7, or 11.8 to version 12.2 without needing to provision a new management node virtual machine.



The Element 12.2 management node is an optional upgrade. It is not required for existing deployments.

What you'll need

- The management node you are intending to upgrade is version 11.3, 11.5, 11.7, or 11.8 and uses IPv4 networking. The management node version 12.2 does not support IPv6.



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 using NetApp Hybrid Cloud Control (HCC). You can access HCC from the following IP: <code><a href="https://<ManagementNodeIP>" class="bare">https://<ManagementNodeIP>;</code>
- If you are updating your management node to version 12.2, you need management services 2.14.60 or later to proceed.
- You have configured an additional network adapter (if required) using the instructions for [configuring an additional storage NIC](#).



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.

- Storage nodes are running Element 11.3 or later.

Steps

1. Configure the management node VM RAM:
 - a. Power off the management node VM.
 - b. Change the RAM of the management node VM from 12GB to 24GB RAM.
 - c. Power on the management node VM.
2. Log in to the management node virtual machine using SSH or console access.
3. Download the [management node ISO](#) for NetApp HCI from the NetApp Support Site to the management node virtual machine.



The name of the ISO is similar to `solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso`

4. Check the integrity of the download by running `md5sum` on the downloaded file and compare 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-<Element release>-patchX-  
XX.X.X.XXXX.iso
```

5. 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-<Element release>-patchX-XX.X.X.XXXX.iso>  
/mnt
```

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

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

```
sudo umount /mnt
```

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

```
sudo rm <path to iso>/solidfire-fdva-<Element release>-patchX-  
XX.X.X.XXXX.iso
```

8. On the 11.3, 11.5, 11.7, or 11.8 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/rtdi/bin/sfrtdi_inplace  
file:///upgrade/casper/filesystem.squashfs sf_upgrade=1
```

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

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



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>
```



If you had previously disabled SSH functionality on the management node, you need to [disable SSH again](#) on the recovered management node. SSH capability that provides [NetApp Support remote support tunnel \(RST\) session access](#) is enabled on the management node by default.

Upgrade a management node to version 12.2 from 11.1 or 11.0

You can perform an in-place upgrade of the management node from 11.0 or 11.1 to version 12.2 without needing to provision a new management node virtual machine.

What you'll need

- Storage nodes are running Element 11.3 or later.



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 version 12.2 does not support IPv6.



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



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.

Steps

- Configure the management node VM RAM:
 - Power off the management node VM.
 - Change the RAM of the management node VM from 12GB to 24GB RAM.
 - Power on the management node VM.
- Log in to the management node virtual machine using SSH or console access.
- Download the [management node ISO](#) for NetApp HCI from the NetApp Support Site to the management node virtual machine.



The name of the ISO is similar to `solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso`

- Check the integrity of the download by running `md5sum` on the downloaded file and compare 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-<Element release>-patchX-XX.X.X.XXXX.iso
```

5. 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-<Element release>-patchX-XX.X.X.XXXX.iso /mnt
```

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

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

```
sudo umount /mnt
```

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

```
sudo rm <path to iso>/solidfire-fdva-<Element release>-patchX-XX.X.X.XXXX.iso
```

8. 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.

- a. 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"
```

- b. 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"
```

- c. 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.

9. On the 12.2 management node, run the `upgrade-mnode` script to retain previous configuration settings.



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.

- a. 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>
```

- b. 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>
```

- c. 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>
```

- d. For multiple storage clusters managed by an existing management node 11.0 or 11.1 with no persistent volumes (the `-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>
```

10. (For all NetApp HCI installations with NetApp Element Plug-in for vCenter Server) Update the vCenter Plug-in on the 12.2 management node by following the steps in the [Upgrade the Element Plug-in for vCenter Server](#) topic.

11. Locate the asset ID for your installation using the management node API:

- a. From a browser, log into the management node REST API UI:
 - i. Go to the storage MVIP and log in.
This action causes certificate to be accepted for the next step.
- b. Open the inventory service REST API UI on the management node:

`https://<ManagementNodeIP>/inventory/1/`

- c. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client`.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the window.
- d. From the REST API UI, select **GET /installations**.
- e. Select **Try it out**.
- f. Select **Execute**.
- g. From the code 200 response body, copy the `id` for the installation.

Your installation has a base asset configuration that was created during installation or upgrade.

12. Locate the hardware tag for your compute node in vSphere:
 - a. Select the host in the vSphere Web Client navigator.
 - b. Select the **Monitor** tab, and select **Hardware Health**.
 - c. The node BIOS manufacturer and model number are listed. Copy and save the value for `tag` for use in a later step.
13. Add a vCenter controller asset for HCI monitoring and Hybrid Cloud Control to the management node known assets:
 - a. Select **POST /assets/{asset_id}/controllers** to add a controller sub-asset.
 - b. Select **Try it out**.
 - c. Enter the parent base asset ID you copied to your clipboard in the `asset_id` field.
 - d. Enter the required payload values with type vCenter and vCenter credentials.
 - e. Select **Execute**.
14. Add a compute node asset to the management node known assets:
 - a. Select **POST /assets/{asset_id}/compute-nodes** to add a compute node sub-asset with credentials for the compute node asset.
 - b. Select **Try it out**.
 - c. Enter the parent base asset ID you copied to your clipboard in the `asset_id` field.
 - d. In the payload, enter the required payload values as defined in the Model tab. Enter `ESXi Host` as `type` and paste the hardware tag you saved during a previous step for `hardware_tag`.
 - e. Select **Execute**.

Migrating from management node version 10.x to 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).
3. Use the TUI to create a new administrator ID and assign a password.
4. In the management node TUI, log in to the management node with the new ID and password and validate that it works.
5. 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
```

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



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**.

9. Log out of vSphere.
10. In a web browser, open the management node registration utility and select **QoSSIOC Service Management**:

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

11. Set the new QoSSIOC password.



The default password is `solidfire`. This password is required to set the new password.

12. Select the **vCenter Plug-in Registration** tab.
13. Select **Update Plug-in**.
14. Enter required values. When you are finished, select **UPDATE**.
15. Log in to vSphere and select **NetApp Element Configuration > mNode Settings**.
16. Select **Actions > Configure**.
17. 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.

18. From the management node registration utility (<https://<mNode 11.1 IP address>:9443>), restart the SIOC service from **QoSSIOC Service Management**.
19. Wait for one minute and check the **NetApp Element Configuration > mNode Settings** tab. This should display the mNode status as **UP**.

If the 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-----
```

20. 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.
21. (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-nma21
```

22. 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.
23. 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
```

24. Verify that ONTAP Select is working by viewing the logs with the following command:

```
journalctl -f | grep -i ots
```

25. Configure Active IQ 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.

26. In vSphere, the **NetApp Element Configuration > mNode Settings** tab should display the mNode status as **UP**.
27. Verify NMA is reporting system alerts and ONTAP Select alerts.
28. If everything is working as expected, shut down and delete management node 10.x VM.

Reconfigure authentication using the management node REST API

You can keep your existing management node if you have sequentially upgraded (1) management services and (2) Element storage. If you have followed a different upgrade order, see the procedures for in-place management node upgrades.

What you'll need

- You have updated your management services to 2.10.29 or later.
- Your storage cluster is running Element 12.0 or later.
- Your management node is 11.3 or later.
- You have sequentially updated your management services followed by upgrading your Element storage.
You cannot reconfigure authentication using this procedure unless you have completed upgrades in the sequence described.

Steps

1. Open the management node REST API UI on the management node:

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

2. Select **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. Select **Authorize** to begin a session.
3. From the REST API UI, select **POST /services/reconfigure-auth**.
4. Select **Try it out**.

5. For the **load_images** parameter, select `true`.

6. Select **Execute**.

The response body indicates that reconfiguration was successful.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Upgrade the Element Plug-in for vCenter Server

For existing vSphere environments with a registered NetApp Element Plug-in for vCenter Server, you can update your plug-in registration after you first update the management services package that contains the plug-in service.

You can update the plug-in registration on vCenter Server Virtual Appliance (vCSA) or Windows using the registration utility. You must change your registration for the vCenter Plug-in on every vCenter Server where you need to use the plug-in.

This upgrade procedure covers the following upgrade scenarios:

- You are upgrading to Element Plug-in for vCenter Server 4.10, 4.9, 4.8, 4.7, 4.6, 4.5, or 4.4.
- You are upgrading to a 7.0, 6.7, or 6.5 HTML5 vSphere Web Client.



The plug-in is not compatible with VMware vCenter Server 6.5 for Element Plug-in for VMware vCenter Server 4.6, 4.7, and 4.8.

- You are upgrading to a 6.7 Flash vSphere Web Client.



The plug-in is not compatible with version 6.7 U2 build 13007421 of the HTML5 vSphere Web Client and other 6.7 U2 builds released prior to update 2a (build 13643870). For more information about supported vSphere versions, see the release notes for [your version of the plug-in](#).

What you'll need

- **Admin privileges:** You have vCenter Administrator role privileges to install a plug-in.
- **vSphere upgrades:** You have performed any required vCenter upgrades before upgrading the NetApp Element Plug-in for vCenter Server. This procedure assumes that vCenter upgrades have already been completed.
- **vCenter Server:** Your vCenter Plug-in version 4.x is registered with a vCenter Server. From the registration utility (<https://<ManagementNodeIP>:9443>), select **Registration Status**, complete the necessary fields, and select **Check Status** to verify that the vCenter Plug-in is already registered and the version number of the current installation.
- **Management services updates:** You have updated your [management services bundle](#) to the latest version. Updates to the vCenter plug-in are distributed using management services updates that are released outside of major product releases for NetApp HCI and SolidFire all-flash storage.
- **Management node upgrades:** You are running a management node that has been [upgraded](#) to version 11.3 or later. vCenter Plug-in 4.4 or later requires a 11.3 or later management node with a modular architecture that provides individual services. Your management node must be powered on with its IP

address or DHCP address configured.

- **Element storage upgrades:** You have a cluster running NetApp Element software 11.3 or later.
- **vSphere Web Client:** You have logged out of the vSphere Web Client before beginning any plug-in upgrade. The web client will not recognize updates made during this process to your plug-in if you do not log out.

Steps

1. Enter the IP address for your management node in a browser, including the TCP port for registration:
<https://<ManagementNodeIP>:9443>
The registration utility UI opens to the **Manage QoSSIOC Service Credentials** page for the plug-in.

The screenshot shows a web-based management interface for the NetApp Element Plug-in. The top navigation bar includes the NetApp logo and the title 'Element Plug-in for vCenter Server Management Node'. Below the title, there are two tabs: 'QoSSIOC Service Management' (which is active) and 'vCenter Plug-in Registration'. The main content area is titled 'Manage QoSSIOC Service Credentials'. On the left, a sidebar menu lists 'QoSSIOC Management', 'Manage Credentials' (which is selected and highlighted in blue), and 'Restart QoSSIOC Service'. The main form contains three password fields: 'Old Password', 'New Password', and 'Confirm Password'. Each field has a placeholder and a descriptive tooltip. Below the fields is a 'SUBMIT CHANGES' button. At the bottom of the page, a footer bar contains the text 'Contact NetApp Support at <http://mysupport.netapp.com>'.

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

Manage vCenter Plug-in

- [Register Plug-in](#)
- [Update Plug-in](#)
- [Unregister Plug-in](#)
- [Registration Status](#)

vCenter Plug-in - Registration

Register version of the NetApp Element Plug-in for vCenter Server with your vCenter server. The Plug-in will not be deployed until a fresh vCenter login after registration.

vCenter Address
Enter the IPv4, IPv6 or DNS name of the vCenter server to register plug-in on.

vCenter User Name
Ensure this user is a vCenter user that has administrative privileges for registration.

vCenter Password
The password for the vCenter user name entered.

Customize URL
Select to customize the Zip file URL.

Plug-in Zip URL
URL of XML initialization file.

REGISTER

Contact NetApp Support at <http://mysupport.netapp.com>

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

4. Confirm or update the following information:

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



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

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



You can select **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 Element Plug-in for vCenter Server documentation about modifying vCenter properties for an in-house (dark site) HTTP server.

5. Select **Update**.

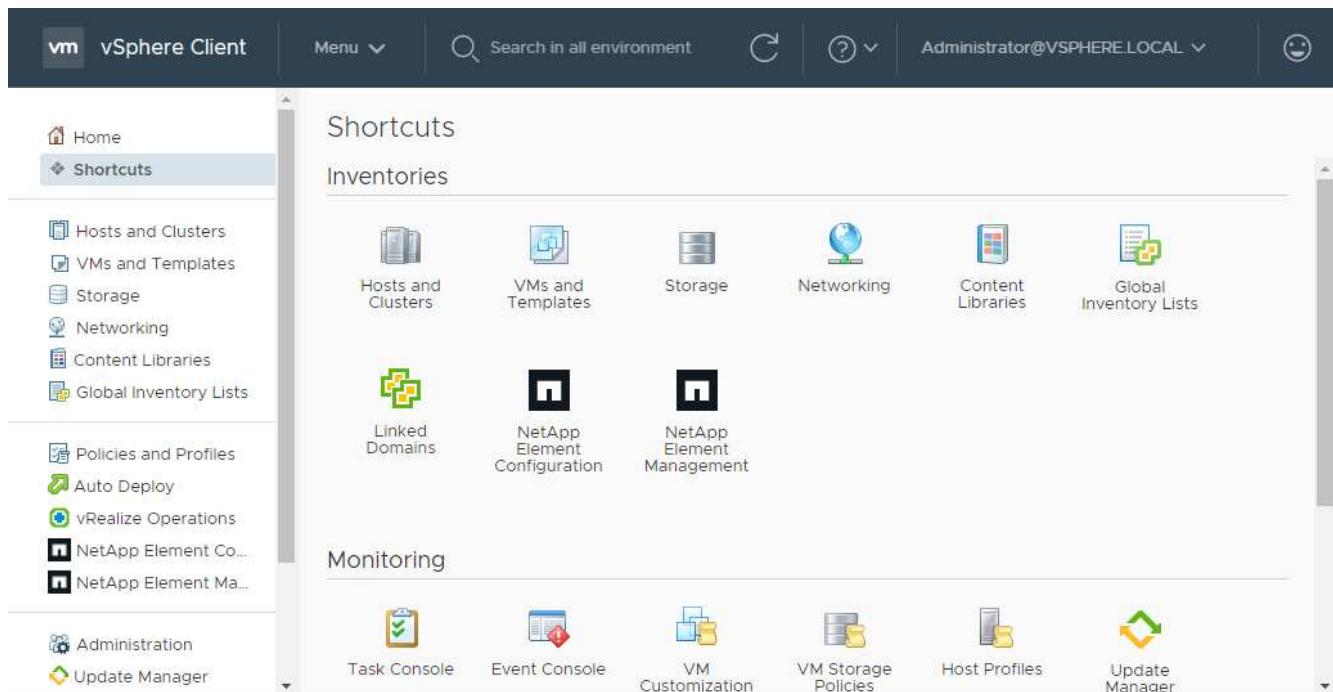
A banner appears in the registration utility UI when the registration is successful.

6. Log in to the vSphere Web Client as a vCenter Administrator. If you are already logged in to the vSphere Web Client, you must first log out and then log in again.



This action creates a new database and completes the installation in the vSphere Web Client.

7. In the vSphere Web Client, look for the following completed tasks in the task monitor to ensure installation has completed: **Download plug-in** and **Deploy plug-in**.
8. Verify that the NetApp Element Configuration and Management extension points appear in the **Shortcuts** tab of the vSphere Web Client and in the side panel.



If the vCenter Plug-in icons are not visible, see [Element Plug-in for vCenter Server](#) documentation about troubleshooting the plug-in.



After upgrading to NetApp Element Plug-in for vCenter Server 4.8 or later with VMware vCenter Server 6.7U1, if the storage clusters are not listed or a server error appears in the **Clusters** and **QoSIOC Settings** sections of the NetApp Element Configuration, see [Element Plug-in for vCenter Server](#) documentation about troubleshooting these errors.

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

You should see the following version details or details of a more recent version:

```
NetApp Element Plug-in Version: 4.10
NetApp Element Plug-in Build Number: 12
```



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.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Run compute node health checks prior to upgrading compute firmware

You must run health checks prior to upgrading compute firmware to ensure all compute nodes in your cluster are ready to be upgraded. Compute node health checks can only be run against compute clusters of one or more managed NetApp HCI compute nodes.

What you'll need

- **Management services:** You have updated to the latest management services bundle (2.11 or later).
- **Management node:** You are running management node 11.3 or later.
- **Element software:** Your storage cluster is running NetApp Element software 11.3 or later.
- **End User License Agreement (EULA):** Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to run compute node health checks:

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Health check options

You can run health checks using NetApp Hybrid Cloud Control (HCC) UI or HCC API:

- [Use NetApp Hybrid Cloud Control to run compute node health checks prior to upgrading firmware](#) (Preferred method)
- [Use API to run compute node health checks prior to upgrading firmware](#)

You can also find out more about compute node health checks that are run by the service:

- [Compute node health checks made by the service](#)

Use NetApp Hybrid Cloud Control to run compute node health checks prior to upgrading firmware

Using NetApp Hybrid Cloud Control (HCC), you can verify that a compute node is ready for a firmware upgrade.

 If you have multiple two-node storage cluster configurations, each within their own vCenter, Witness Nodes health checks might not report accurately. Therefore, when you are ready to upgrade ESXi hosts, you must only shut down the Witness Node on the ESXi host that is being upgraded. You must ensure that you always have one Witness Node running in your NetApp HCI installation by powering off the Witness Nodes in an alternate fashion.

Steps

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select the **Compute firmware** tab.
5.  Select the health check for the cluster you want to check for upgrade readiness.
6. On the **Compute Health Check** page, select **Run Health Check**.
7. If there are issues, the page provides a report. Do the following:
 - a. Go to the specific KB article listed for each issue or perform the specified remedy.
 - b. If a KB is specified, complete the process described in the relevant KB article.
 - c. After you have resolved cluster issues, select **Re-Run Health Check**.

After the health check completes without errors, the compute nodes in the cluster are ready to upgrade. See [Update compute node firmware](#) to proceed.

Use API to run compute node health checks prior to upgrading firmware

You can use REST API to verify that compute nodes in a cluster are ready to be upgraded. The health check verifies that there are no obstacles to upgrading, such as ESXi host issues or other vSphere issues. You will need to run compute node health checks for each compute cluster in your environment.

Steps

1. Locate the controller ID and cluster ID:
 - a. Open the inventory service REST API UI on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```
 - b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnnode-client` if the value is not already populated.
 - iii. Select **Authorize** to begin a session.
 - c. From the REST API UI, select **GET /installations**.
 - d. Select **Try it out**.
 - e. Select **Execute**.
 - f. From the code 200 response body, copy the "`id`" for the installation you plan to use for health checks.
 - g. From the REST API UI, select **GET /installations/{id}**.
 - h. Select **Try it out**.
 - i. Enter the installation ID.
 - j. Select **Execute**.

k. From the code 200 response body, copy the IDs for each of the following:

- i. The cluster ID ("clusterID")
- ii. A controller ID ("controllerId")

```
{  
  "_links": {  
    "collection":  
      "href": "https://10.117.187.199/inventory/1/installations",  
    "self":  
      "href": "https://10.117.187.199/inventory/1/installations/xx94f6f0-12a6-  
412f-8b5e-4cf2z58329x0"  
  },  
  "compute": {  
    "errors": [],  
    "inventory": {  
      "clusters": [  
        {  
          "clusterId": "domain-1",  
          "controllerId": "abc12c3a-aa87-4e33-9f94-xx588c2cdcf6",  
          "datacenterName": "NetApp-HCI-Datacenter-01",  
          "installationId": "xx94f6f0-12a6-412f-8b5e-  
4cf2z58329x0",  
          "installationName": "test-nde-mnode",  
          "inventoryType": "managed",  
          "name": "NetApp-HCI-Cluster-01",  
          "summary": {  
            "nodeCount": 2,  
            "virtualMachineCount": 2  
          }  
        }  
      ]  
    }  
  }  
}
```

2. Run health checks on the compute nodes in the cluster:

a. Open the compute service REST API UI on the management node:

```
https://<ManagementNodeIP>/vcenter/1/
```

b. Select **Authorize** and complete the following:

- i. Enter the cluster user name and password.
- ii. Enter the client ID as `mnode-client` if the value is not already populated.
- iii. Select **Authorize** to begin a session.

c. Select **POST /compute/{CONTROLLER_ID}/health-checks**.

- d. Select **Try it out**.
- e. Enter the "controllerId" you copied from the previous step in the **Controller_ID** parameter field.
- f. In the payload, enter the "clusterId" that you copied from the previous step as the "cluster" value and remove the "nodes" parameter.

```
{
  "cluster": "domain-1"
}
```

- g. Select **Execute** to run a health check on the cluster.

The code 200 response gives a "resourceLink" URL with the task ID appended that is needed to confirm the health check results.

```
{
  "resourceLink":
  "https://10.117.150.84/vcenter/1/compute/tasks/ [This is the task ID
  for health check task results]",
  "serviceName": "vcenter-v2-svc",
  "taskId": "ab12c345-06f7-42d7-b87c-7x64x56x321x",
  "taskName": "VCenter service health checks"
}
```

- h. Copy the task ID portion of the "resourceLink" URL to verify the task result.
3. Verify the result of the health checks:
 - a. Return to the compute service REST API UI on the management node:

<https://<ManagementNodeIP>/vcenter/1/>
 - b. Select **GET /compute/tasks/{task_id}**.
 - c. Select **Try it out**.
 - d. Enter the task ID portion of the "resourceLink" URL from the **POST /compute/{CONTROLLER_ID}/health-checks** code 200 response in the `task_id` parameter field.
 - e. Select **Execute**.
 - f. If the `status` returned indicates that there were problems regarding compute node health, do the following:
 - i. Go to the specific KB article (KbLink) listed for each issue or perform the specified remedy.
 - ii. If a KB is specified, complete the process described in the relevant KB article.
 - iii. After you have resolved cluster issues, run **POST /compute/{CONTROLLER_ID}/health-checks** again (see step 2).

If health checks complete without issues, the response code 200 indicates a successful result.

Compute node health checks made by the service

Compute health checks, whether performed by HCC or API methods, make the following checks per node. Depending on your environment, some of these checks might be skipped. You should re-run health checks after resolving any detected issues.

Check description	Node/cluster	Action needed to resolve	Knowledgebase article with procedure
Is DRS enabled and fully automated?	Cluster	Turn on DRS and make sure it is fully automated.	See this KB . NOTE: If you have standard licensing, put the ESXi host into maintenance mode and ignore this health check failure warning.
Is DPM disabled in vSphere?	Cluster	Turn off Distributed Power Management.	See this KB .
Is HA admission control disabled in vSphere?	Cluster	Turn off HA admission control.	See this KB .
Is FT enabled for a VM on a host in the cluster?	Node	Suspend Fault Tolerance on any affected virtual machines.	See this KB .
Are there critical alarms in vCenter for the cluster?	Cluster	Launch vSphere and resolve and/or acknowledge any alerts before proceeding.	No KB needed to resolve issue.
Are there generic/global informational alerts in vCenter?	Cluster	Launch vSphere and resolve and/or acknowledge any alerts before proceeding.	No KB needed to resolve issue.
Are management services up to date?	HCI system	You must update management services before you perform an upgrade or run pre-upgrade health checks.	No KB needed to resolve issue. See this article for more information.
Are there errors on the current ESXi node in vSphere?	Node	Launch vSphere and resolve and/or acknowledge any alerts before proceeding.	No KB needed to resolve issue.
Is virtual media mounted to a VM on a host in the cluster?	Node	Unmount all virtual media disks (CD/DVD/floppy) from the VMs.	No KB needed to resolve issue.
Is BMC version the minimum required version that has RedFish support?	Node	Manually update your BMC firmware.	No KB needed to resolve issue.
Is ESXi host up and running?	Node	Start your ESXi host.	No KB needed to resolve issue.

Check description	Node/cluster	Action needed to resolve	Knowledgebase article with procedure
Do any virtual machines reside on local ESXi storage?	Node/VM	Remove or migrate local storage attached to virtual machines.	No KB needed to resolve issue.
Is BMC up and running?	Node	Power on your BMC and ensure it is connected to a network this management node can reach.	No KB needed to resolve issue.
Are there partner ESXi host(s) available?	Node	Make one or more ESXi host(s) in cluster available (not in maintenance mode) to migrate virtual machines.	No KB needed to resolve issue.
Are you able to connect with BMC via IPMI protocol?	Node	Enable IPMI protocol on Baseboard Management Controller (BMC).	No KB needed to resolve issue.
Is ESXi host mapped to hardware host (BMC) correctly?	Node	The ESXi host is not mapped to the Baseboard Management Controller (BMC) correctly. Correct the mapping between ESXi host and hardware host.	No KB needed to resolve issue. See this article for more information.
What is the status of the Witness Nodes in the cluster? None of the witness nodes identified are up and running.	Node	A Witness Node is not running on an alternate ESXi host. Power on the Witness Node on an alternate ESXi host and re-run the health check. One Witness Node must be running in the HCI installation at all times.	See this KB
What is the status of the Witness Nodes in the cluster? The witness node is up and running on this ESXi host and the alternate witness node is not up and running.	Node	A Witness Node is not running on an alternate ESXi host. Power on the Witness Node on an alternate ESXi host. When you are ready to upgrade this ESXi host, shut down the witness node running on this ESXi host and re-run the health check. One Witness Node must be running in the HCI installation at all times.	See this KB

Check description	Node/cluster	Action needed to resolve	Knowledgebase article with procedure
What is the status of the Witness Nodes in the cluster? Witness node is up and running on this ESXi host and the alternate node is up but is running on the same ESXi host.	Node	Both Witness Nodes are running on this ESXi host. Relocate one Witness Node to an alternate ESXi host. When you are ready to upgrade this ESXi host, shut down the Witness Node remaining on this ESXi host and re-run the health check. One Witness Node must be running in the HCI installation at all times.	See this KB
What is the status of the Witness Nodes in the cluster? Witness node is up and running on this ESXi host and the alternate witness node is up and running on another ESXi host.	Node	A Witness Node is running locally on this ESXi host. When you are ready to upgrade this ESXi host, shut down the Witness Node only on this ESXi host and re-run the health check. One Witness Node must be running in the HCI installation at all times.	See this KB

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Update compute node drivers

For any H-Series compute node, you can update the drivers used on the nodes using VMware Update Manager.

What you'll need

See the firmware and driver matrix for your hardware at [supported firmware and ESXi driver versions](#).

About this task

Perform only one of these update operations at a time.

Steps

1. Browse to the [NetApp HCI software downloads](#) page and select the download link for correct version of NetApp HCI.
2. Select **ESXI_drivers** from the drop-down list.
3. Accept the End User License Agreement.
4. Download the driver package for your node type and ESXi version.
5. Extract the downloaded driver bundle on your local computer.



The NetApp driver bundle includes one or more VMware Offline Bundle ZIP files; do not extract these ZIP files.

6. After upgrading the firmware on the compute nodes, go to **VMware Update Manager** in VMware vCenter.
7. Import the driver offline bundle file for the compute nodes into the **Patch Repository**.
 - For VMware ESXi 7.0, all the necessary drivers for NetApp H610C, H615C, H410C, and Hx00E compute nodes and their build-in system components are included in the standard VMware ESXi 7.0 installation ISO image. You do not require additional or updated drivers for NetApp HCI compute nodes running VMware ESXi 7.0 (and updates).
 - For VMware ESXi 6.x, perform the following steps to import the driver offline bundle file:
 - a. Select the **Updates** tab.
 - b. Select **UPLOAD FROM FILE**.
 - c. Browse to the offline bundle that was previously downloaded and select **IMPORT**.
8. Create a new host baseline for the compute node.
9. Choose **Host Extension** for Name and Type and select all imported driver packages to be included in the new baseline.
10. In the **Host and Clusters** menu in vCenter, select the cluster with the compute nodes you would like to update and navigate to the **Update Manager** tab.
11. Select **Remediate** and the select the newly created host baseline. Ensure that drivers included in the baseline are selected.
12. Proceed through the wizard to the **Host Remediation Options** and ensure that the **Do Not Change VM Power State** option is selected to keep virtual machines online during the driver update.



If VMware Distributed Resource Scheduler (DRS) is enabled on the cluster (this is the default in NetApp HCI installations), virtual machines will automatically be migrated to other nodes in the cluster.

13. Proceed to the **Ready to Complete** page in the wizard and select **Finish**.

The drivers for all compute nodes in the cluster are updated one node at a time while virtual machines stay online.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Upgrade compute node firmware

For H-series compute nodes, you can upgrade the firmware for hardware components such as the BMC, BIOS, and NIC. To upgrade compute node firmware, you can use the NetApp Hybrid Cloud Control UI, REST API, a USB drive with the latest firmware image, or the BMC UI.

After the upgrade, the compute node boots into ESXi and works as before, retaining the configuration.

What you'll need

- **Compute drivers:** You have upgraded your compute node drivers. If compute node drivers are not compatible with the new firmware, the upgrade will not start. See the [Interoperability Matrix Tool \(IMT\)](#) for driver and firmware compatibility information, and check the latest [compute node firmware release notes](#) for important late-breaking firmware and driver details.
- **Admin privileges:** You have cluster administrator and BMC administrator permissions to perform the upgrade.
- **System ports:** If you are using NetApp Hybrid Cloud Control for upgrades, you have ensured that the necessary ports are open. See [Network ports](#) for more information.
- **Minimum BMC and BIOS versions:** The node you intend to upgrade using NetApp Hybrid Cloud Control meets the following minimum requirements:

Model	Minimum BMC version	Minimum BIOS version
H410C	All versions supported (no upgrade required)	All versions supported (no upgrade required)
H610C	3.96.07	3B01
H615C	4.68.07	3B08.CO



H615C compute nodes must update BMC firmware to version 4.68 using the [compute firmware bundle 2.27](#) to enable NetApp Hybrid Cloud Control to perform future firmware upgrades.



For a complete matrix of firmware and driver firmware for your hardware, see [supported firmware and ESXi driver versions](#).

- **BIOS boot order:** Manually change the boot order in the BIOS setup for each node to ensure USB CD/DVD appears in the boot list. See this [article](#) for more information.
- **BMC credentials:** Update the credentials NetApp Hybrid Cloud Control uses to connect to the compute node BMC. You can do this using either the NetApp Hybrid Cloud Control [UI](#) or [API](#). Updating BMC information prior to upgrade refreshes the inventory and ensures that management node services are aware of all hardware parameters needed to complete the upgrade.
- **Attached media:** Disconnect any physical USB or ISO before starting a compute node upgrade.
- **KVM ESXi console:** Close all open Serial-Over-LAN (SOL) sessions and active KVM sessions in the BMC UI before starting a compute node upgrade.
- **Witness Node requirements:** In two- and three-node storage clusters, one [Witness Node](#) must be running in the NetApp HCI installation at all times.
- **Compute node health check:** You have verified that the node is ready to be upgraded. See [Run compute node health checks prior to upgrading compute firmware](#).
- **End User License Agreement (EULA):** Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to upgrade compute node firmware:

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

About this task

In production environments, upgrade the firmware on one compute node at a time.



The ESXi host must be taken out of lockdown mode prior to running a health check and proceeding with the firmware upgrade. See [How to disable lockdown mode on ESXi host](#) and [VMware lockdown mode behavior](#) for more information.

For NetApp Hybrid Cloud Control UI or API upgrades, your ESXi host will be automatically placed in maintenance mode during the upgrade process if you have the DRS feature and required licensing. The node will be rebooted and after the upgrade process is complete, the ESXi host will be taken out of maintenance mode. For USB and BMC UI options, you will need to place the ESXi host in maintenance mode manually, as described in each procedure.

Upgrade options

Choose the option that is relevant to your upgrade scenario:

- [Use NetApp Hybrid Cloud Control UI to upgrade a compute node](#) (Recommended)
- [Use NetApp Hybrid Cloud Control API to upgrade a compute node](#)
- [Use a USB drive imaged with the latest compute firmware bundle](#)
- [Use the Baseboard Management Controller \(BMC\) user interface \(UI\)](#)

Use NetApp Hybrid Cloud Control UI to upgrade a compute node

Starting with management services 2.14, you can upgrade a compute node using the NetApp Hybrid Cloud Control UI. From the list of nodes, you must select the node to upgrade. The **Current Versions** tab shows the current firmware versions and the **Proposed Versions** tab shows the available upgrade versions, if any.



For a successful upgrade, ensure that the health check on the vSphere cluster is successful.



Upgrading the NIC, BIOS, and BMC can take approximately 60 minutes per node depending on the speed of network connectivity between the management node and the BMC host.



Using the NetApp Hybrid Cloud Control UI to upgrade compute firmware on H300E/H500E/H700E compute nodes is no longer supported. To upgrade, it is recommended that you use a [USB drive](#) or the [BMC UI](#) to mount the compute firmware bundle.

What you'll need

- If your management node is not connected to the internet, you have downloaded the compute firmware bundle from the [NetApp Support Site](#).



You should extract the `TAR.GZ` file to a `TAR` file, and then extract the `TAR` file to the compute firmware bundle.

Steps

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select **Compute firmware**.
5. Select the cluster you are upgrading.

You will see the nodes in the cluster listed along with the current firmware versions and newer versions, if available for upgrade.

6. Select **Browse** to upload the compute firmware bundle that you downloaded from the [NetApp Support Site](#).
7. Wait for the upload to complete. A progress bar shows the status of the upload.



The file upload will happen in the background if you navigate away from the browser window.

An on-screen message is displayed after the file is successfully uploaded and validated. Validation might take several minutes.

8. Select the compute firmware bundle.
9. Select **Begin Upgrade**.

After you select **Begin Upgrade**, the window shows failed health checks, if any.



The upgrade cannot be paused after you begin. Firmware will be updated sequentially in the following order: NIC, BIOS, and BMC. Do not log in to the BMC UI during upgrade. Logging into the BMC terminates the Hybrid Cloud Control Serial-Over-LAN (SOL) session that monitors upgrade process.

10. If the health checks at the cluster or node level passed with warnings, but without critical failures, you will see **Ready to be Upgraded**. Select **Upgrade Node**.



While the upgrade is in progress, you can leave the page and come back to it later to continue monitoring the progress. During the upgrade, the UI shows various messages about the status of the upgrade.



While upgrading the firmware on H610C compute nodes, do not open the Serial-Over-LAN (SOL) console through the BMC web UI. This might cause the upgrade to fail.

The UI displays a message after the upgrade is complete. You can download logs after the upgrade is complete. For information about the various upgrade status changes, see [Upgrade status changes](#).



If a failure happens during the upgrade, NetApp Hybrid Cloud Control will reboot the node, take it out of maintenance mode, and display the failure status with a link to the error log. You can download the error log, which contains specific instructions or links to KB articles, to diagnose and correct any issue. For additional insight into compute node firmware upgrade issues using NetApp Hybrid Cloud Control, see this [KB article](#).

Upgrade status changes

Here are the different states that the UI shows before, during, and after the upgrade process:

Upgrade state	Description
Node failed one or more health checks. Expand to view details.	One or more health checks failed.
Error	An error has occurred during the upgrade. You can download the error log and send it to NetApp Support.
Unable to Detect	This status is displayed if NetApp Hybrid Cloud Control is unable to query the compute node when the compute node asset does not have the hardware tag.
Ready to be Upgraded.	All the health checks passed successfully, and the node is ready to be upgraded.
An error has occurred during the upgrade.	The upgrade fails with this notification when a critical error occurs. Download the logs by selecting the Download Logs link to help resolve the error. You can try upgrading again after you resolve the error.
Node upgrade is in progress.	The upgrade is in progress. A progress bar shows the upgrade status.

Use NetApp Hybrid Cloud Control API to upgrade a compute node

You can use APIs to upgrade each compute node in a cluster to the latest firmware version. You can use an automation tool of your choice to run the APIs. The API workflow documented here uses the REST API UI available on the management node as an example.



Using the NetApp Hybrid Cloud Control UI to upgrade compute firmware on H300E/H500E/H700E compute nodes is no longer supported. To upgrade, it is recommended that you use a [USB drive](#) or the [BMC UI](#) to mount the compute firmware bundle.

What you'll need

Compute node assets, including vCenter and hardware assets, must be known to management node assets. You can use the inventory service APIs to verify assets (<https://<ManagementNodeIP>/inventory/1/>).

Steps

1. Go to the NetApp HCI software [downloads page](#) and download the latest compute firmware bundle to a device that is accessible to the management node.



For dark site upgrades, you can reduce upload time if the upgrade package and the management node are both local.

2. Upload the compute firmware bundle to the management node:

a. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/package-repository/1/
```

b. Select **Authorize** and complete the following:

- i. Enter the cluster user name and password.
- ii. Enter the client ID as `mnode-client`.
- iii. Select **Authorize** to begin a session.
- iv. Close the authorization window.

c. From the REST API UI, select **POST /packages**.

d. Select **Try it out**.

e. Select **Browse** and select the compute firmware bundle.

f. Select **Execute** to initiate the upload.

g. From the response, copy and save the compute firmware bundle ID ("id") for use in a later step.

3. Verify the status of the upload.

a. From the REST API UI, select **GET /packages/{id}/status**.

b. Select **Try it out**.

c. Enter the compute firmware bundle ID you copied in the previous step in **id**.

d. Select **Execute** to initiate the status request.

The response indicates `state` as `SUCCESS` when complete.

e. From the response, copy and save the compute firmware bundle name ("name") and version ("version") for use in a later step.

4. Locate the compute controller ID and node hardware ID for the node you intend to upgrade:

a. Open the inventory service REST API UI on the management node:

```
https://<ManagementNodeIP>/inventory/1/
```

b. Select **Authorize** and complete the following:

- i. Enter the cluster user name and password.
- ii. Enter the client ID as `mnode-client`.
- iii. Select **Authorize** to begin a session.
- iv. Close the authorization window.

c. From the REST API UI, select **GET /installations**.

d. Select **Try it out**.

e. Select **Execute**.

- f. From the response, copy the installation asset ID ("id").
- g. From the REST API UI, select **GET /installations/{id}**.
- h. Select **Try it out**.
- i. Paste the installation asset ID into the **id** field.
- j. Select **Execute**.
- k. From the response, copy and save the cluster controller ID ("controllerId") and node hardware ID ("hardwareId") for use in a later step:

```

"compute": {
  "errors": [],
  "inventory": {
    "clusters": [
      {
        "clusterId": "Test-1B",
        "controllerId": "a1b23456-c1d2-11e1-1234-a12bcdef123a",

```

```

"nodes": [
  {
    "bmcDetails": {
      "bmcAddress": "10.111.0.111",
      "credentialsAvailable": true,
      "credentialsValidated": true
    },
    "chassisSerialNumber": "111930011231",
    "chassisSlot": "D",
    "hardwareId": "123a4567-01b1-1243-a12b-11ab11ab0a15",
    "hardwareTag": "00000000-0000-0000-0000-ab1c2de34f5g",
    "id": "e1111d10-1a1a-12d7-1a23-ab1cde23456f",
    "model": "H410C",

```

- 5. Run the compute node firmware upgrade:

- a. Open the hardware service REST API UI on the management node:

```
https://<ManagementNodeIP>/hardware/2/
```

- b. Select **Authorize** and complete the following:
 - i. Enter the cluster user name and password.
 - ii. Enter the client ID as `mnode-client`.
 - iii. Select **Authorize** to begin a session.
 - iv. Close the authorization window.

- c. Select **POST /nodes/{hardware_id}/upgrades**.
- d. Select **Try it out**.
- e. Enter the hardware host asset ID ("hardwareId" saved from a previous step) in the parameter field.
- f. Do the following with the payload values:
 - i. Retain the values "force": false and "maintenanceMode": true so that health checks are performed on the node and the ESXi host is set to maintenance mode.
 - ii. Enter the cluster controller ID ("controllerId" saved from a previous step).
 - iii. Enter the compute firmware bundle name package version that you saved from a previous step.

```
{
  "config": {
    "force": false,
    "maintenanceMode": true
  },
  "controllerId": "a1b23456-c1d2-11e1-1234-a12bcdef123a",
  "packageName": "compute-firmware-12.2.109",
  "packageVersion": "12.2.109"
}
```

- g. Select **Execute** to initiate the upgrade.



The upgrade cannot be paused after you begin. Firmware will be updated sequentially in the following order: NIC, BIOS, and BMC. Do not log in to the BMC UI during upgrade. Logging into the BMC terminates the Hybrid Cloud Control Serial-Over-LAN (SOL) session that monitors upgrade process.

- h. Copy the upgrade task ID that is part of the resource link ("resourceLink") URL in the response.
6. Verify the upgrade progress and results:
- a. Select **GET /task/{task_id}/logs**.
 - b. Select **Try it out**.
 - c. Enter the task ID from the previous step in **task_Id**.
 - d. Select **Execute**.
 - e. Do one of the following if there are problems or special requirements during the upgrade:

Option	Steps
You need to correct cluster health issues due to <code>failedHealthChecks</code> message in the response body.	<ol style="list-style-type: none"> 1. Go to the specific KB article listed for each issue or perform the specified remedy. 2. If a KB is specified, complete the process described in the relevant KB article. 3. After you have resolved cluster issues, reauthenticate if needed and select POST /nodes/{hardware_id}/upgrades. 4. Repeat the steps as described previously in the upgrade step.
The upgrade fails and the mitigation steps are not listed in upgrade log.	<ol style="list-style-type: none"> 1. See this KB article (login required).

- f. Run the **GET /task/{task_id}/logs** API multiple times, as needed, until the process is complete.

During the upgrade, the `status` indicates `running` if no errors are encountered. As each step finishes, the `status` value changes to `completed`.

The upgrade has finished successfully when the `status` for each step is `completed` and the `percentageCompleted` value is 100.

7. (Optional) Confirm upgraded firmware versions for each component:

- a. Open the hardware service REST API UI on the management node:

`https://<ManagementNodeIP>/hardware/2/`

- b. Select **Authorize** and complete the following:

- Enter the cluster user name and password.
- Enter the client ID as `mnode-client`.
- Select **Authorize** to begin a session.
- Close the authorization window.

- c. From the REST API UI, select **GET /nodes/{hardware_id}/upgrades**.

- (Optional) Enter date and status parameters to filter the results.
- Enter the hardware host asset ID ("hardwareId" saved from a previous step) in the parameter field.
- Select **Try it out**.
- Select **Execute**.
- Verify in the response that firmware for all components has been successfully upgraded from the previous version to the latest firmware.

Use a USB drive imaged with the latest compute firmware bundle

You can insert a USB drive with the latest compute firmware bundle downloaded to a USB port on the compute node. As an alternative to using the USB thumb drive method described in this procedure, you can mount the

compute firmware bundle on the compute node using the **Virtual CD/DVD** option in the Virtual Console in the Baseboard Management Controller (BMC) interface. The BMC method takes considerably longer than the USB thumb drive method. Ensure that your workstation or server has the necessary network bandwidth and that your browser session with the BMC does not time out.

What you'll need

- If your management node is not connected to the internet, you have downloaded the compute firmware bundle from the [NetApp Support Site](#).



You should extract the `TAR.GZ` file to a `TAR` file, and then extract the `TAR` file to the compute firmware bundle.

Steps

1. Use the Etcher utility to flash the compute firmware bundle to a USB drive.
 2. Place the compute node in maintenance mode using VMware vCenter, and evacuate all virtual machines from the host.
-
- If VMware Distributed Resource Scheduler (DRS) is enabled on the cluster (this is the default in NetApp HCI installations), virtual machines will automatically be migrated to other nodes in the cluster.
3. Insert the USB thumb drive into a USB port on the compute node and reboot the compute node using VMware vCenter.
 4. During the compute node POST cycle, press **F11** to open the Boot Manager. You may need to press **F11** multiple times in quick succession. You can perform this operation by connecting a video/keyboard or by using the console in BMC.
 5. Select **One Shot > USB Flash Drive** from the menu that appears. If the USB thumb drive does not appear in the menu, verify that USB Flash Drive is part of the legacy boot order in the BIOS of the system.
 6. Press **Enter** to boot the system from the USB thumb drive. The firmware flash process begins.

After firmware flashing is complete and the node reboots, it might take a few minutes for ESXi to start.

7. After the reboot is complete, exit maintenance mode on the upgraded compute node using vCenter.
8. Remove the USB flash drive from the upgraded compute node.
9. Repeat this task for other compute nodes in your ESXi cluster until all compute nodes are upgraded.

Use the Baseboard Management Controller (BMC) user interface (UI)

You must perform the sequential steps to load the compute firmware bundle and reboot the node to the compute firmware bundle to ensure that the upgrade is successful. The compute firmware bundle should be located on the system or virtual machine (VM) hosting the web browser. Verify that you have downloaded the compute firmware bundle before you start the process.



The recommendation is to have the system or VM and the node on the same network.



It takes approximately 25 to 30 minutes for the upgrade via the BMC UI.

- [Upgrade firmware on H410C and H300E/H500E/H700E nodes](#)

- [Upgrade firmware on H610C/H615C nodes](#)

Upgrade firmware on H410C and H300E/H500E/H700E nodes

If your node is part of a cluster, you must place the node in maintenance mode before the upgrade, and take it out of maintenance mode after the upgrade.



Ignore the following informational message you see during the process: Untrusty Debug Firmware Key is used, SecureFlash is currently in Debug Mode

Steps

1. If your node is part of a cluster, place it in maintenance mode as follows. If not, skip to step 2.
 - a. Log in to the VMware vCenter web client.
 - b. Right-click the host (compute node) name and select **Maintenance Mode > Enter Maintenance Mode**.
 - c. Select **OK**.
VMs on the host will be migrated to another available host. VM migration can take time depending on the number of VMs that need to be migrated.



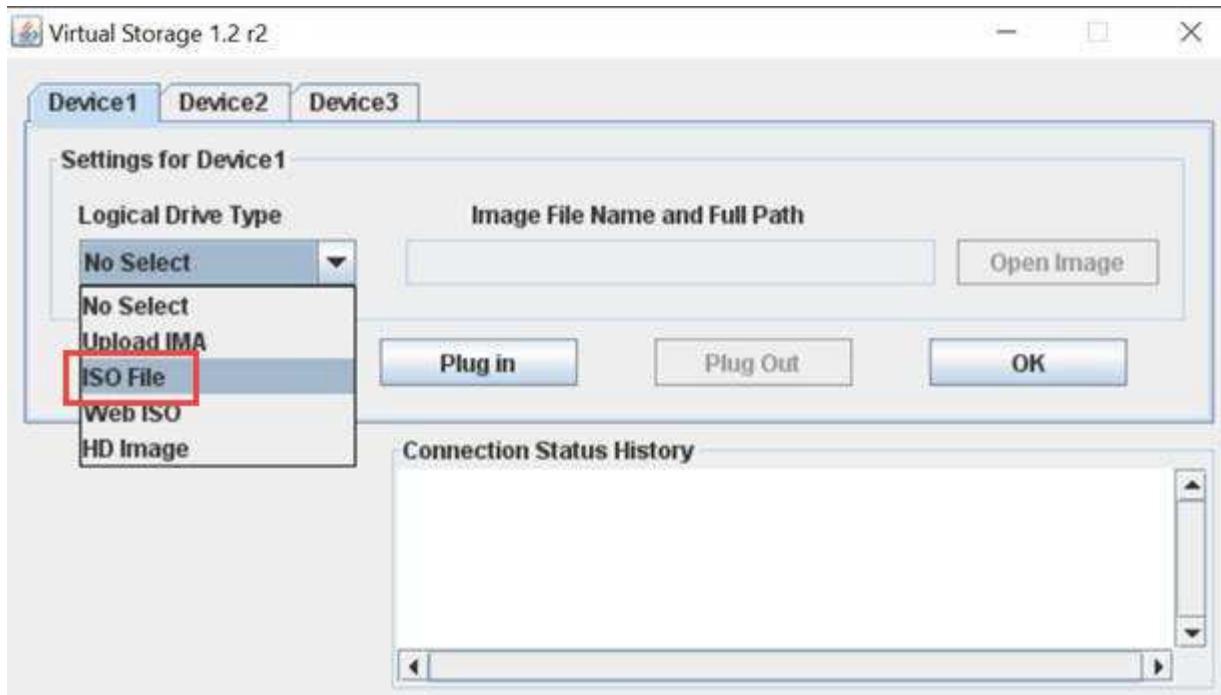
Ensure that all the VMs on the host are migrated before you proceed.

2. Navigate to the BMC UI, <https://BMCIP/#login>, where BMCIP is the IP address of the BMC.
3. Log in using your credentials.
4. Select **Remote Control > Console Redirection**.
5. Select **Launch Console**.



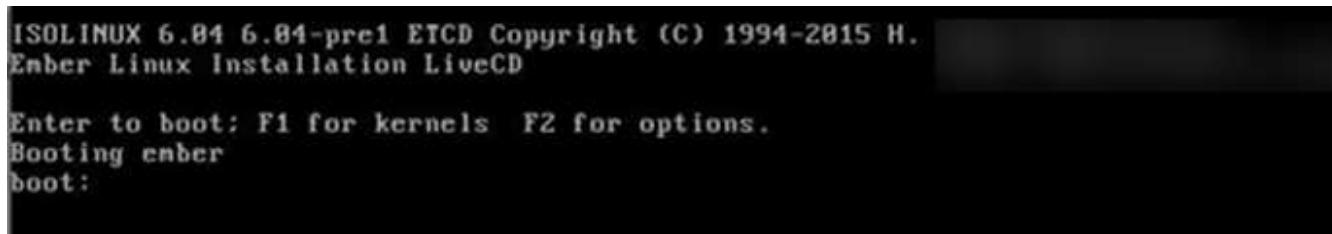
You might have to install Java or update it.

6. When the console opens, select **Virtual Media > Virtual Storage**.
7. On the **Virtual Storage** screen, select **Logical Drive Type**, and select **ISO File**.



8. Select **Open Image** to browse to the folder where you downloaded the compute firmware bundle file, and select the compute firmware bundle file.
9. Select **Plug In**.
10. When the connection status shows Device#: VM Plug-in OK!!, select **OK**.
11. Reboot the node by pressing **F12** and selecting **Restart** or selecting **Power Control > Set Power Reset**.
12. During reboot, press **F11** to select the boot options and load the compute firmware bundle. You might have to press F11 a few times before the boot menu is displayed.

You will see the following screen:



13. On the above screen, press **Enter**. Depending on your network, it might take a few minutes after you press **Enter** for the upgrade to begin.



Some of the firmware upgrades might cause the console to disconnect and/or cause your session on the BMC to disconnect. You can log back into the BMC, however some services, such as the console, may not be available due to the firmware upgrades. After the upgrades have completed, the node will perform a cold reboot, which can take approximately five minutes.

14. Log back in to the BMC UI and select **System** to verify the BIOS version and build time after booting to the OS. If the upgrade completed correctly, you see the new BIOS and BMC versions.



The BIOS version will not show the upgraded version until the node has finished fully booting.

15. If the node is part of a cluster, complete the steps below. If it is a standalone node, no further action is needed.
 - a. Log in to the VMware vCenter web client.
 - b. Take the host out of maintenance mode. This might show a disconnected red flag. Wait until all statuses are cleared.
 - c. Power on any of the remaining VMs that were powered off.

Upgrade firmware on H610C/H615C nodes

The steps vary depending on whether the node is standalone or part of a cluster. The procedure can take approximately 25 minutes and includes powering the node off, uploading the compute firmware bundle, flashing the devices, and powering the node back on after the upgrade.

Steps

1. If your node is part of a cluster, place it in maintenance mode as follows. If not, skip to step 2.
 - a. Log in to the VMware vCenter web client.
 - b. Right-click the host (compute node) name and select **Maintenance Mode > Enter Maintenance Mode**.
 - c. Select **OK**.
VMs on the host will be migrated to another available host. VM migration can take time depending on the number of VMs that need to be migrated.



Ensure that all the VMs on the host are migrated before you proceed.

2. Navigate to the BMC UI, <https://BMCIP/#login>, where BMC IP is the IP address of the BMC.
3. Log in using your credentials.
4. Select **Remote Control > Launch KVM (Java)**.
5. In the console window, select **Media > Virtual Media Wizard**.



6. Select **Browse** and select the compute firmware .iso file.
7. Select **Connect**.
A popup indicating success is displayed, along with the path and device showing at the bottom. You can close the **Virtual Media** window.



8. Reboot the node by pressing **F12** and selecting **Restart** or selecting **Power Control > Set Power Reset**.
9. During reboot, press **F11** to select the boot options and load the compute firmware bundle.
10. Select **AMI Virtual CDROM** from the list displayed and select **Enter**. If you do not see AMI Virtual CDROM in the list, go into the BIOS and enable it in the boot list. The node will reboot after you save. During the reboot, press **F11**.



11. On the screen displayed, select **Enter**.



Some of the firmware upgrades might cause the console to disconnect and/or cause your session on the BMC to disconnect. You can log back into the BMC, however some services, such as the console, might not be available due to the firmware upgrades. After the upgrades have completed, the node will perform a cold reboot, which can take approximately five minutes.

12. If you get disconnected from the console, select **Remote Control** and select **Launch KVM** or **Launch KVM (Java)** to reconnect and verify when the node has finished booting back up. You might need multiple reconnects to verify that the node booted successfully.



During the powering on process, for approximately five minutes, the KVM console displays **No Signal**.

13. After the node is powered on, select **Dashboard > Device Information > More info** to verify the BIOS and BMC versions. The upgraded BIOS and BMC versions are displayed. The upgraded version of the BIOS will not be displayed until the node has fully booted up.
14. If you placed the node in maintenance mode, after the node boots to ESXi, right-click the host (compute node) name, and select **Maintenance Mode > Exit Maintenance Mode**, and migrate the VMs back to the host.
15. In vCenter, with the host name selected, configure and verify the BIOS version.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Upgrade your vSphere components for a NetApp HCI system with the Element Plug-in for vCenter Server

When you upgrade the VMware vSphere components of your NetApp HCI installation, there are some additional steps you will need to take for the Element Plug-in for vCenter Server.

Steps

1. For vCSA upgrades, [clear](#) QoSSIOC settings in the plug-in (**NetApp Element Configuration > QoSSIOC Settings**). The **QoSSIOC Status** field displays **Not Configured** after the process is complete.
2. For vCSA and Windows upgrades, [unregister](#) the plug-in from the vCenter Server with which it is associated using the registration utility.
3. [Upgrade vSphere, including vCenter Server, ESXi, VMs, and other VMware components.](#)



When you upgrade to VMware vCenter Server 7.0 U3, the Element Plug-in fails to deploy. To resolve this issue using Spring Framework 4, see [this KB article](#).



When upgrading ESXi for compute nodes for a [two-node cluster](#), upgrade only one compute node at a time so that only one witness node is temporarily unavailable and cluster quorum can be maintained.

4. [Register](#) the Element Plug-in for vCenter Server again with vCenter.
5. [Add clusters](#) using the plug-in.
6. [Configure QoSSIOC settings](#) using the plug-in.
7. [Enable QoSSIOC](#) for all datastores controlled by the plug-in.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Two-Node Storage Cluster Technical Report](#)

Expand your NetApp HCI system

Expansion overview

You can expand your NetApp HCI system by using NetApp Hybrid Cloud Control. You can expand storage or compute resources separately or expand them at the same time.



New and spare H610S storage nodes might have additional installation requirements based on the existing Element software version of the storage cluster. Contact NetApp Support for more information.

After installing the node in the NetApp HCI chassis, you use NetApp Hybrid Cloud Control to configure NetApp HCI to use the new resources. NetApp HCI detects the existing network configuration and offers you configuration options within the existing networks and VLANs, if any.



If you recently expanded your installation and the new assets were not added automatically to your configuration, you might need to add the assets manually.

See [Management node overview](#).

NetApp HCI uses VMware Enhanced vMotion Compatibility (EVC) to ensure vMotion functionality when there are compute nodes with different CPU generations in the vSphere cluster. When EVC is required for expansion, NetApp HCI enables it automatically whenever possible.

In the following situations, you might need to manually change EVC settings in the vSphere client to complete expansion:

- The existing compute nodes have a newer CPU generation than the compute nodes you are trying to add.
- The controlling vCenter instance does not support the required EVC level.
- The compute nodes you are trying to add have an older CPU generation than the EVC setting of the controlling vCenter instance.



When expanding NetApp HCI compute or storage resources in the NetApp Deployment Engine, you should connect to the vCenter instance that manages your existing NetApp HCI compute nodes.

Find more information

- [Expand NetApp HCI compute resources](#)
- [Expand NetApp HCI storage resources](#)
- [Expand NetApp HCI storage and compute resources at the same time](#)
- [NetApp Element Plug-in for vCenter Server](#)

Expand NetApp HCI storage resources

After you finish NetApp HCI deployment, you can expand and configure NetApp HCI storage resources by using NetApp Hybrid Cloud Control.

What you'll need

- Ensure that you have free and unused IPv4 addresses on the same network segment as existing nodes (each new node must be installed on the same network as existing nodes of its type).
- Ensure that you have one of the following types of SolidFire storage cluster accounts:
 - The native administrator account that was created during initial deployment
 - A custom user account with Cluster Admin, Drives, Volumes, and Nodes permissions
- Ensure that you have performed the following actions with each new node:
 - Installed the new node in the NetApp HCI chassis by following the [installation instructions](#).
 - Cabled and powered on the new node
- Ensure that you have the management IPv4 address of an already installed storage node. You can find the IP address in the **NetApp Element Management > Cluster > Nodes** tab of the NetApp Element Plug-in for vCenter Server.
- Ensure that each new node uses the same network topology and cabling as the existing storage or compute clusters.



When you are expanding storage resources, storage capacity should be split evenly across all chassis for the best reliability.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.

3. Select **Expand** at the top right corner of the interface.

The browser opens the NetApp Deployment Engine.

4. Log in to the NetApp Deployment Engine by providing the local NetApp HCI storage cluster administrator credentials.



You cannot log in using Lightweight Directory Access Protocol credentials.

5. On the **Welcome** page, select **No** and select **Continue**.

6. On the **Available Inventory** page, select the storage nodes you want to add and select **Continue**.

7. On the **Network Settings** page, some of the network information has been detected from the initial deployment. Each new storage node is listed by serial number, and you need to assign the new network information to it. For each new storage node, complete the following steps:

- a. **Hostname**: If NetApp HCI detected a naming prefix, copy it from the Detected Naming Prefix field, and insert it as the prefix for the new unique hostname you add in the Hostname field.
- b. **Management Address**: Enter a management IP address for the new storage node that is within the management network subnet.
- c. **Storage (iSCSI) IP Address**: Enter an iSCSI IP address for the new storage node that is within the iSCSI network subnet.

d. Select **Continue**.



NetApp HCI might take some time to validate the IP addresses you enter. The **Continue** button becomes available when IP address validation completes.

8. On the **Review** page in the Network Settings section, new nodes are shown in the bold text. To make changes in any section, do the following:

a. Select **Edit** for that section.

b. After you finish, select **Continue** on any subsequent pages to return to the Review page.

9. **Optional:** If you do not want to send cluster statistics and support information to NetApp hosted Active IQ servers, clear the final checkbox.

This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve issues before production is impacted.

10. Select **Add Nodes**.

You can monitor the progress while NetApp HCI adds and configures the resources.

11. **Optional:** Verify that any new storage nodes are visible in the Element Plug-in for vCenter Server.



If you expanded a two-node storage cluster to four nodes or more, the pair of Witness Nodes previously used by the storage cluster are still visible as standby virtual machines in vSphere. The newly expanded storage cluster does not use them; if you want to reclaim VM resources, you can [manually remove](#) the Witness Node virtual machines.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Expand NetApp HCI compute resources

After you finish NetApp HCI deployment, you can expand and configure NetApp HCI compute resources by using NetApp Hybrid Cloud Control.

What you'll need

- Ensure that the vSphere instance of NetApp HCI is using vSphere Enterprise Plus licensing if you are expanding a deployment with Virtual Distributed Switches.
- Ensure that none of the vCenter or vSphere instances in use with NetApp HCI have expired licenses.
- Ensure that you have free and unused IPv4 addresses on the same network segment as existing nodes (each new node must be installed on the same network as existing nodes of its type).
- Ensure that you have the vCenter administrator account credentials ready.
- Ensure that you have performed the following actions with each new node:
 - Installed the new node in the NetApp HCI chassis by following the [installation instructions](#).
 - Cabled and powered on the new node
- Ensure that each new node uses the same network topology and cabling as the existing storage or

compute clusters.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.
3. Select **Expand** at the top right corner of the interface.

The browser opens the NetApp Deployment Engine.

4. Log in to the NetApp Deployment Engine by providing the local NetApp HCI storage cluster administrator credentials.



You cannot log in using Lightweight Directory Access Protocol credentials.

5. On the **Welcome** page, select **Yes** and select **Continue**.
6. On the **End User License** page, read the VMware End User License Agreement and select **I accept** to accept the terms and select **Continue**.
7. On the **vCenter** page, complete the following steps:
 - a. Enter a FQDN or IP address and administrator credentials for the vCenter instance associated with your NetApp HCI installation.
 - b. Select **Continue**.
 - c. Select a vSphere datacenter where you want to add the compute nodes, or select **Create New Datacenter** to add the compute nodes to a new datacenter.



If you select **Create New Datacenter**, the Cluster field is automatically populated.

- d. If you selected an existing datacenter, select a vSphere cluster with which the new compute nodes should be associated.



If NetApp HCI cannot recognize the network settings of the cluster you have selected for expansion, ensure that the vmkernel and vmnic mapping for the management, storage and vMotion networks are set to the deployment defaults. See [supported networking changes](#) for more information.

- e. Select **Continue**.
8. On the **ESXi Credentials** page, enter an ESXi root password for the compute node or nodes you are adding.

You should use the same password that was created during the initial NetApp HCI deployment.

9. Select **Continue**.
10. If you created a new vSphere datacenter cluster, on the **Network Topology** page, select a network topology to match the new compute nodes you are adding.



You can select the two-cable option only if your compute nodes are using the two-cable topology and the existing NetApp HCI deployment is configured with VLAN IDs.

11. On the **Available Inventory** page, select the nodes you want to add to the existing NetApp HCI installation.



For some compute nodes, you might need to enable EV at the highest level that your vCenter version supports before you can add them to your installation. You need to use the vSphere client to enable EVC for these compute nodes. After you enable it, refresh the Inventory page and try adding the compute nodes again.

12. Select **Continue**.

13. **Optional:** If you created a new vSphere datacenter cluster, on the **Network Settings** page, import network information from an existing NetApp HCI deployment by selecting the **Copy Setting from an Existing Cluster** checkbox.

This populates the default gateway and subnet information for each network.

14. On the **Network Settings** page, some of the network information has been detected from the initial deployment. Each new compute node is listed by serial number, and you need to assign new network information to it. For each new compute node, complete the following steps:

- a. **Hostname:** If NetApp HCI detected a naming prefix, copy it from the **Detected Naming Prefix** field, and insert it as the prefix for the new hostname.
- b. **Management IP Address:** Enter a management IP address for the new compute node that is within the management network subnet.
- c. **vMotion IP Address:** Enter a vMotion IP address for the new compute node that is within the vMotion network subnet.
- d. **iSCSI A - IP Address:** Enter an IP address for the first iSCSI port of the compute node that is in the iSCSI network subnet.
- e. **iSCSI B - IP Address:** Enter an IP address for the second iSCSI port of the compute node that is in the iSCSI network subnet
- f. Select **Continue**.

15. On the **Review** page in the Network Settings section, new nodes are shown in the bold text. To make changes in any section, do the following:

- a. Select **Edit** for that section.
- b. After you finish, select **Continue** on any subsequent pages to return to the **Review** page.

16. **Optional:** If you do not want to send cluster statistics and support information to NetApp hosted SolidFire Active IQ servers, clear the final checkbox.

This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve issues before production is impacted.

17. Select **Add Nodes**.

You can monitor the progress while NetApp HCI adds and configures the resources.

18. **Optional:** Verify that any new compute nodes are visible in the VMware vSphere Web Client.

Find more information

- [NetApp HCI Compute and Storage Nodes Installation and Setup Instructions](#)
- [VMware Knowledge Base: Enhanced vMotion Compatibility \(EVC\) processor support](#)

Expand NetApp HCI storage and compute resources at the same time

After you finish NetApp HCI deployment, you can expand and configure NetApp HCI storage and compute resources at the same time by using NetApp Hybrid Cloud Control.

What you'll need

- Ensure that the vSphere instance of NetApp HCI is using vSphere Enterprise Plus licensing if you are expanding a deployment with Virtual Distributed Switches.
- Ensure that none of the vCenter or vSphere instances in use with NetApp HCI have expired licenses.
- Ensure that you have the vCenter administrator account credentials ready.
- Ensure that you have free and unused IPv4 addresses on the same network segment as existing nodes (each new node must be installed on the same network as existing nodes of its type).
- Ensure that you have one of the following types of SolidFire storage cluster accounts:
 - The native administrator account that was created during initial deployment
 - A custom user account with Cluster Admin, Drives, Volumes, and Nodes permissions
- Ensure that you have performed the following actions with each new node:
 - Installed the new node in the NetApp HCI chassis by following the [installation instructions](#).
 - Cabled and powered on the new node
- Ensure that you have the management IPv4 address of an already installed storage node. You can find the IP address in the **NetApp Element Management > Cluster > Nodes** tab of the NetApp Element Plug-in for vCenter Server.
- Ensure that each new node uses the same network topology and cabling as the existing storage or compute clusters.

About this task

- You can intermix the H410C compute node with existing NetApp HCI compute and storage nodes in the same chassis and cluster.
- You cannot intermix compute nodes and BPU-enabled compute nodes in the same cluster. If you select a GPU-enabled compute node, CPU-only compute nodes become unselectable, and vice versa.
- If you are adding compute nodes with CPU generations that are different than the CPU generation of the existing compute nodes and Enhanced vMotion Compatibility (EVC) is disabled on the controlling vCenter instance, you must enable EVC before proceeding. This ensures vMotion functionality after expansion is complete.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.

3. Select **Expand** at the top right corner of the interface.

The browser opens the NetApp Deployment Engine.

4. Log in to the NetApp Deployment Engine by providing the local NetApp HCI storage cluster administrator credentials.



You cannot log in using Lightweight Directory Access Protocol credentials.

5. On the **Welcome** page, select **Yes** and select **Continue**.

6. On the **End User License** page, read the VMware End User License Agreement and select **I accept** to accept the terms and select **Continue**.

7. On the **vCenter** page, complete the following steps:

a. Enter a FQDN or IP address and administrator credentials for the vCenter instance associated with your NetApp HCI installation.

b. Select **Continue**.

c. Select a vSphere datacenter where you want to add the compute nodes, or select **Create New Datacenter** to add the compute nodes to a new datacenter.



If you select Create New Datacenter, the Cluster field is automatically populated.

d. If you selected an existing datacenter, select a vSphere cluster with which the new compute nodes should be associated.



If NetApp HCI cannot recognize the network settings of the cluster you have selected for expansion, ensure that the vmkernel and vmnic mapping for the management, storage and vMotion networks are set to the deployment defaults. See [supported networking changes](#) for more information.

e. Select **Continue**.

8. On the **ESXi Credentials** page, enter an ESXi root password for the compute node or nodes you are adding.

You should use the same password that was created during the initial NetApp HCI deployment.

9. Select **Continue**.

10. If you created a new vSphere datacenter cluster, on the **Network Topology** page, select a network topology to match the new compute nodes you are adding.



You can select the two-cable option only if your compute nodes are using the two-cable topology and the existing NetApp HCI deployment is configured with VLAN IDs.

11. On the **Available Inventory** page, select the storage and compute nodes you want to add and select **Continue**.



For some compute nodes, you might need to enable EV at the highest level that your vCenter version supports before you can add them to your installation. You need to use the vSphere client to enable EVC for these compute nodes. After you enable it, refresh the Inventory page and try adding the compute nodes again.

12. Select **Continue**.

13. **Optional:** If you created a new vSphere datacenter cluster, on the **Network Settings** page, import network information from an existing NetApp HCI deployment by selecting the **Copy Setting from an Existing Cluster** checkbox.

This populates the default gateway and subnet information for each network.

14. On the **Network Settings** page, some of the network information has been detected from the initial deployment. Each new storage node is listed by serial number, and you need to assign the new network information to it. For each new storage node, complete the following steps:

- a. **Hostname:** If NetApp HCI detected a naming prefix, copy it from the Detected Naming Prefix field, and insert it as the prefix for the new unique hostname you add in the Hostname field.
- b. **Management Address:** Enter a management IP address for the new storage node that is within the management network subnet.
- c. **Storage (iSCSI) IP Address:** Enter an iSCSI IP address for the new storage node that is within the iSCSI network subnet.
- d. Select **Continue**.



NetApp HCI might take some time to validate the IP addresses you enter. The **Continue** button becomes available when IP address validation completes.

15. On the **Review** page in the Network Settings section, new nodes are shown in the bold text. To make changes in any section, do the following:

- a. Select **Edit** for that section.
- b. After you finish, select **Continue** on any subsequent pages to return to the Review page.

16. **Optional:** If you do not want to send cluster statistics and support information to NetApp hosted Active IQ servers, clear the final checkbox.

This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve issues before production is impacted.

17. Select **Add Nodes**.

You can monitor the progress while NetApp HCI adds and configures the resources.

18. **Optional:** Verify that any new nodes are visible in the VMware vSphere Web Client (for compute nodes) or the Element Plug-in for vCenter Server (for storage nodes).



If you expanded a two-node storage cluster to four nodes or more, the pair of Witness Nodes previously used by the storage cluster are still visible as standby virtual machines in vSphere. The newly expanded storage cluster does not use them; if you want to reclaim VM resources, you can [manually remove](#) the Witness Node virtual machines.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Compute and Storage Nodes Installation and Setup Instructions](#)
- [VMware Knowledge Base: Enhanced vMotion Compatibility \(EVC\) processor support](#)

Remove Witness Nodes after expanding cluster

After you expand a two-node storage cluster to four or more nodes, you can delete the pair of Witness Nodes to free up compute resources in your NetApp HCI installation. The Witness Nodes previously used by the storage cluster are still visible as standby virtual machines (VM) in vSphere Web Client.

About this task

Witness Nodes are not required in clusters with more than four storage nodes. This is an optional procedure if you want to free up CPU and memory after you expand your two-node cluster to four or more nodes.



Verify that no cluster faults or errors are reported. You can find information about system alerts by selecting **Reporting > Alerts** in the NetApp Element Management extension point in vSphere.

Steps

1. From vSphere, access the NetApp Element Management extension point from the **Shortcuts** tab or the side panel.
2. Select **NetApp Element Management > Cluster > Nodes**.

3. Select the checkbox for the Witness Node that you want to delete, and select **Actions > Remove**.
4. Confirm the action in the prompt.

5. Select **Hosts and Clusters**.
6. Navigate to the Witness Node VM that you removed earlier.
7. Right-click the VM and power it off.
8. Right-click the VM that you powered off, and select **Delete from Disk**.
9. Confirm the action in the prompt.

Find more information

- [NetApp HCI Two-Node Storage Cluster | TR-4823](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Use Rancher on NetApp HCI

Rancher on NetApp HCI overview

Rancher is a complete software stack for teams adopting containers. Rancher addresses the operational and security challenges of managing multiple Kubernetes clusters across different infrastructures, while providing DevOps teams with integrated tools for running containerized workloads.

Deploying Rancher on NetApp HCI deploys the Rancher control plane, also referred to as the *Rancher server*, and enables you to create on-premises Kubernetes clusters. You deploy the Rancher control plane by using the NetApp Hybrid Cloud Control.

After deployment, using the Rancher control Plane, you provision, manage, and monitor Kubernetes clusters used by Dev and Ops teams. Dev and Ops teams can use Rancher to perform activities on user clusters that reside on NetApp HCI itself, a public cloud provider, or any other infrastructure that Rancher enables.

Benefits of Rancher on NetApp HCI

- Ease of installation: You do not need to learn how to install and configure Rancher. You can deploy a template-based implementation, which was jointly developed by NetApp HCI and Rancher.
- Lifecycle management: In a manual Rancher implementation, updates for the Rancher server application or the Rancher Kubernetes Engine (RKE) cluster are not automated. Rancher on NetApp HCI provides the ability for updates to the management cluster, that includes the Rancher server and the RKE.

What you can do with Rancher on NetApp HCI

With Rancher on NetApp HCI, you can:

- Deploy services across cloud providers and your private cloud.
- Port the apps and data across a hybrid cloud architecture regardless of cloud location without compromising service-level agreements.
- Spin up cloud-native applications yourself.
- Centralize management of multiple clusters (new and existing).
- Perform orchestration of hybrid cloud Kubernetes-based applications.

Technical Support option

Using Rancher on NetApp HCI and Kubernetes open-source software includes free deployment and usage. License keys are not required.

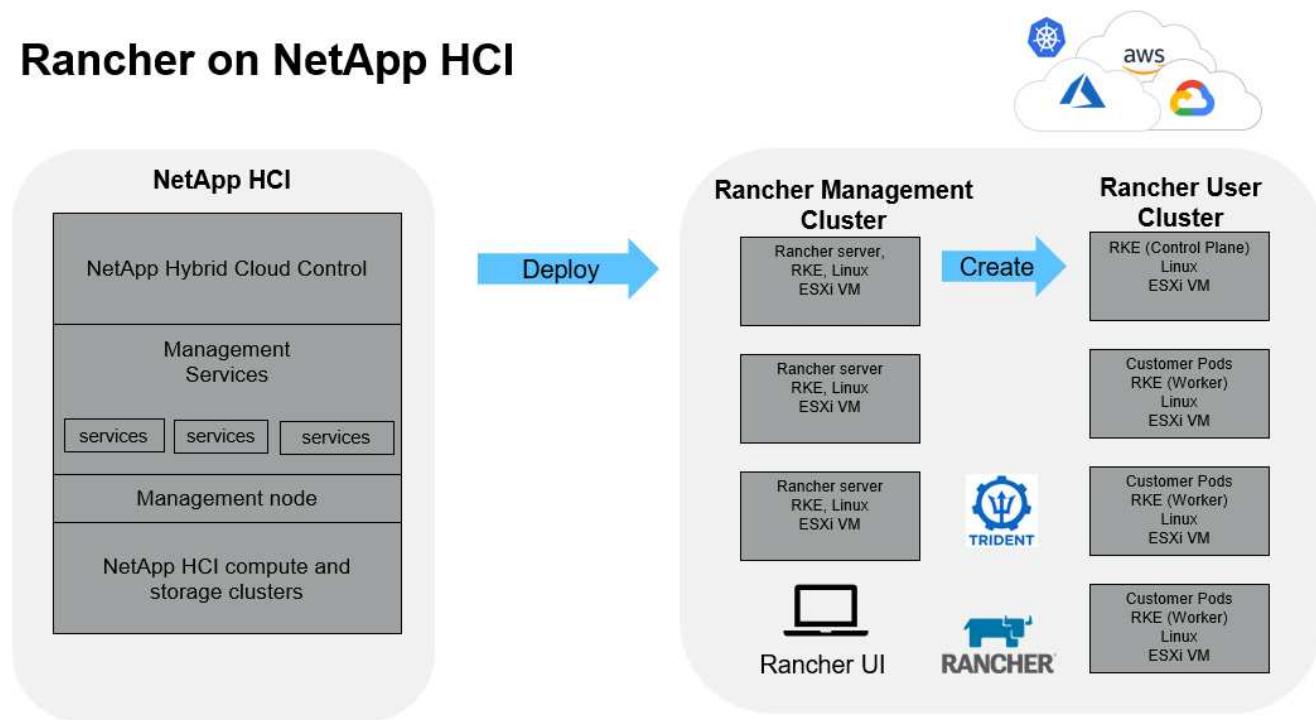
You can choose a NetApp Rancher Support option to obtain core-based, Rancher enterprise support.



Rancher Support is not included in your NetApp Support Edge agreement. Contact NetApp Sales or your reseller for options. If you purchase Rancher Support from NetApp, you will receive an email with instructions.

Rancher on NetApp HCI architecture and components

Here is an overview of the various components of Rancher on NetApp HCI:



- **NetApp Hybrid Cloud Control:** This interface enables you to deploy Rancher on NetApp HCI and NetApp Element software, required for Rancher on NetApp HCI.



You can use NetApp Hybrid Cloud Control also to upgrade management services, expand your system, collect logs, and monitor your installation.

- **Management services:** Management services run on the management node and enable you to deploy Rancher on NetApp HCI using NetApp Hybrid Cloud Control.
- **Management cluster:** Rancher on NetApp HCI deploys three virtual machines on the Rancher management cluster, which you can see using NetApp Hybrid Cloud Control, vCenter Server, or the Rancher user interface. The management cluster virtual machines host the Rancher server, the Rancher Kubernetes Engine (RKE), and the Linux OS.



For the best performance and greater security, consider using a dedicated Kubernetes cluster for the Rancher management server. You should not run your user workloads on the management cluster.

- **User clusters:** The downstream Kubernetes user clusters run your apps and services. Any cluster that you deploy from Rancher or import into Rancher is a user cluster.
- **Trident:** A Trident catalog is available to Rancher on NetApp HCI and runs in the user clusters. Inclusion of this catalog simplifies the Trident deployment to user clusters.

Find more information

- [Rancher documentation about architecture](#)

- [NetApp HCI Resources page](#)

Rancher on NetApp HCI concepts

Learn basic concepts related to Rancher on NetApp HCI.

- **Rancher server or Control plane:** The Rancher control plane, sometimes called the *Rancher Server*, provisions, manages, and monitors Kubernetes clusters used by Development and Operations teams.
- **Catalogs:** Catalogs are GitHub repositories or Helm Chart repositories filled with applications that are ready-made for deployment. Rancher provides the ability to use a catalog of Helm charts that make it easy to deploy applications repeatedly. Rancher includes two types of catalogs: built-in global catalogs and custom catalogs. Trident is deployed as a catalog. See [Rancher documentation about catalogs](#).
- **Management cluster:** Rancher on NetApp HCI deploys three virtual machines on the Rancher management cluster, which you can see using Rancher, Hybrid Cloud Control, and the vCenter Plug-in. The management cluster virtual machines host the Rancher server, the Rancher Kubernetes Engine (RKE), and the Linux OS.
- **User clusters:** These downstream Kubernetes clusters run your apps and services. In Kubernetes installations of Rancher, the management cluster should be separate from the user clusters. Any cluster that a Rancher user deploys from Rancher, or imports into Rancher, is considered a user cluster.
- **Rancher node template:** Hybrid Cloud Control uses a Rancher node template to make deployment simpler.

See [Rancher documentation about node templates](#).

Trident software and persistent storage concepts

Trident, itself a Kubernetes-native application, runs directly within a Kubernetes cluster. With Trident, Kubernetes users (such as developers, data scientists, and Kubernetes administrators) can create, manage, and interact with persistent storage volumes in the standard Kubernetes format that they are already familiar with. With Trident, NetApp solutions can meet persistent volume claims that are made by Kubernetes clusters.

With Rancher, you can use a persistent volume, one that exists independently of any specific pod and with its own lifetime. Using Trident to manage persistent volume claims (PVCs) insulates the developers creating pods from the lower-level implementation details of the storage that they are accessing.

When a containerized application issues a persistent volume claim (PVC) request, Trident dynamically provisions storage per the parameters requested against the NetApp Element software storage layer in NetApp HCI.

A Trident catalog is available to Rancher on NetApp HCI and runs in the user clusters. As part of the Rancher on NetApp HCI implementation, a Trident installer is available in the Rancher catalog by default. Inclusion of this catalog simplifies the Trident deployment to user clusters.

See [Install Trident with Rancher on NetApp HCI](#).

For details, visit the [Trident documentation](#).

Find more information

- [Rancher documentation about architecture](#)

- [Kubernetes terminology for Rancher](#)
- [NetApp HCI Resources page](#)

Requirements for Rancher on NetApp HCI

Before you install Rancher on NetApp HCI, ensure your environment and your NetApp HCI system meet these requirements.

 If you accidentally deploy Rancher on NetApp HCI with incorrect information (such as an incorrect Rancher server FQDN), there is no way to correct the deployment without removing it and redeploying. You will need to remove the Rancher on NetApp HCI instance and then redeploy Rancher on NetApp HCI from NetApp Hybrid Cloud Control UI. See [Remove a Rancher installation on NetApp HCI](#) for more information.

Node requirements

- Ensure that your NetApp HCI system has at least three compute nodes; this is required for full resiliency. Rancher on NetApp HCI is not supported on storage-only configurations.
- Ensure that the datastore you intend to use for the Rancher on NetApp HCI deployment has at least 60GB of free space.
- Ensure that your NetApp HCI cluster is running management services version 2.17 or later.

Node details

Rancher on NetApp HCI deploys a three-node management cluster.

All nodes have the following characteristics:

vCPU	RAM (GB)	Disk (GB)
2	8	20

Network requirements

- Ensure that the network that you intend to deploy the Rancher on NetApp HCI management cluster has a route to the management node management network.
- Rancher on NetApp HCI supports DHCP addresses for the control plane (Rancher server) and user clusters, but we recommend static IP addresses for production environments. Ensure that you have allocated the necessary static IP addresses if you are deploying in a production environment.
 - Rancher server requires three static IP addresses.
 - Each user cluster requires as many static IP addresses as nodes in the cluster. For example, a user cluster with four nodes requires four static IP addresses.
 - If you plan on using DHCP addressing for the Rancher control plane or user clusters, ensure that the DHCP lease duration is at least 24 hours.
- If you need to use an HTTP proxy to enable internet access for Rancher on NetApp HCI, you need to make a pre-deployment change to the management node. Log in to your management node using SSH and follow the [instructions](#) in the Docker documentation to manually update the proxy settings for Docker.
- If you enable and configure a proxy server during deployment, the following IP address ranges and

domains are automatically added to the Rancher server noProxy settings:

```
127.0.0.0/8, 10.0.0.0/8, 172.16.0.0/12, 192.168.0.0/16, .svc,  
.cluster.local
```

- Ensure that your management node can use DNS to resolve the host name <any IP address>.nip.io to an IP address. This is the DNS provider used during deployment; if the management node cannot resolve this URL, deployment will fail.
- Ensure that you have set up DNS records for each static IP address you need.

VMware vSphere requirements

- Ensure that the VMware vSphere instance you are using is version 6.5, 6.7, or 7.0.
- You can use a vSphere Standard Switch (VSS) networking configuration, but if you do, ensure that the virtual switches and physical hosts used for Rancher VMs can access all the same port groups, in the same way that you would ensure for regular VMs.

Deployment considerations

You might want to review the following considerations:

- Types of deployments
 - Demo deployments
 - Production deployments
- Rancher FQDN

 Rancher on NetApp HCI is not resilient to node failures unless you configure some type of network load balancing. As a simple solution, create a round robin DNS entry for the three static IP addresses reserved for Rancher server. These DNS entries should resolve to the Rancher server FQDN that you will use to access the Rancher server host, which serves the Rancher web UI once deployment is complete.

Types of deployments

You can deploy Rancher on NetApp HCI in the following ways:

- **Demo deployments:** If DHCP is available in the targeted deployment environment and you want to demo the Rancher on NetApp HCI capability, then a DHCP deployment makes the most sense.

In this deployment model, the Rancher UI is accessible from each of the three nodes in the management cluster.

If your organization does not use DHCP, you can still try it out by using four static IP addresses allocated prior the deployment, similar to what you would do for a production deployment.

- **Production deployments:** For production deployments or when DHCP is not available in the targeted deployment environment, a little more pre-deployment work is required. The first step is to obtain three consecutive IP addresses. You enter the first during the deployment.

We recommend using L4 load balancing or round-robin DNS configuration for production environments.

This requires a fourth IP address and separate entry in your DNS configuration.

- **L4 load balancing:** This is a technique where a virtual machine or container hosting an application like nginx is configured to distribute requests among the three nodes of the management cluster.
- **Round-robin DNS:** This is a technique where a single host name is configured in the DNS system that rotates requests among the three hosts that form the management cluster.

Rancher FQDN

The installation requires assignment of a Rancher URL, which includes the fully qualified domain name (FQDN) of the host where the Rancher UI will be served after the installation is complete.

In all cases the Rancher UI is accessible in your browser over https protocol (port 443).

Production deployments require an FQDN configured that load balances across the management cluster nodes. Without using FQDN and load balancing, the environment is not resilient and is suitable only for demo environments.

Required ports

Ensure that the list of ports in the "Ports for Rancher Server Nodes on RKE" section of the **Rancher Nodes** section of the official [Rancher documentation](#) are open in your firewall configuration to and from the nodes running Rancher server.

Required URLs

The following URLs should be accessible from the hosts where the Rancher control plane resides:

URL	Description
https://charts.jetstack.io/	Kubernetes integration
https://releases.rancher.com/server-charts/stable	Rancher software downloads
https://entropy.ubuntu.com/	Ubuntu entropy service for random number generation
https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo/v1.3.1/install.sh	VMware guest additions
https://download.docker.com/linux/ubuntu/gpg	Docker Ubuntu GPG public key
https://download.docker.com/linux/ubuntu	Docker download link
https://hub.docker.com/	Docker Hub for NetApp Hybrid Cloud Control

Deploy Rancher on NetApp HCI

To use Rancher on your NetApp HCI environment, you first deploy Rancher on NetApp HCI.



Before starting the deployment, be sure to check the datastore free space and other [requirements for Rancher on NetApp HCI](#).



Rancher Support is not included in your NetApp Support Edge agreement. Contact NetApp Sales or your reseller for options. If you purchase Rancher Support from NetApp, you will receive an email with instructions.

What happens when you deploy Rancher on NetApp HCI?

The deployment involves the following steps, each described further:

- Use the NetApp Hybrid Cloud Control to initiate the deployment.
- The Rancher deployment creates a management cluster, which includes three virtual machines.

Each virtual machine is assigned all Kubernetes roles for both the Control Plane and Worker. This means that the Rancher UI is available on each node.

- The Rancher Control Plane (or *Rancher Server*) is also installed, using the NetApp HCI node template in Rancher for easier deployment. The Rancher Control Plane automatically works with the configuration used in the NetApp Deployment Engine, which was used to build the NetApp HCI infrastructure.
- After deployment, you will receive an email from NetApp providing you with the option to register for NetApp Support on Rancher deployments on NetApp HCI.
- After deployment, Dev and Ops teams can then deploy their user clusters, similar to any Rancher deployment.

Steps to deploy Rancher on NetApp HCI

- [Access the NetApp Hybrid Cloud Control](#)
- [Deploy Rancher on NetApp HCI](#)
- [Verify your deployment by using vCenter Server](#)

Access the NetApp Hybrid Cloud Control

To begin the deployment, access the NetApp Hybrid Cloud Control.

1. Open the IP address of the management node in a web browser. For example:

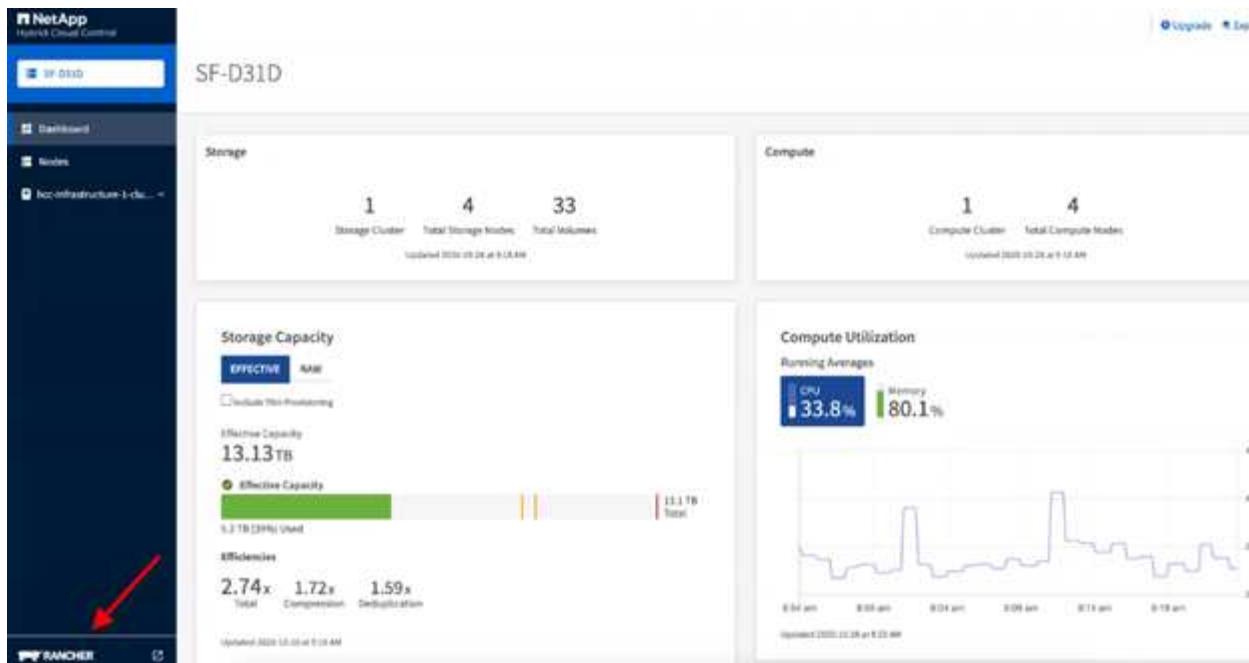
```
https://<ManagementNodeIP>
```

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.

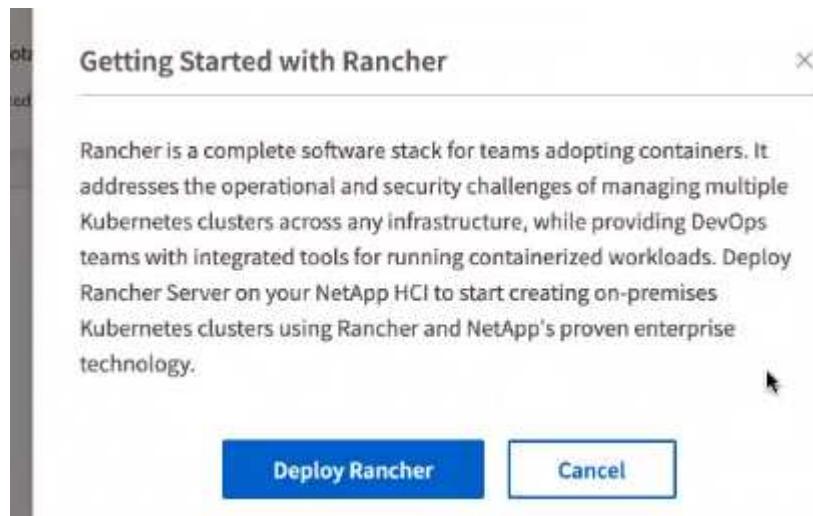
The NetApp Hybrid Cloud Control interface appears.

Deploy Rancher on NetApp HCI

1. From the Hybrid Cloud Control, select the **Rancher** icon in the lower left of the navigation bar.

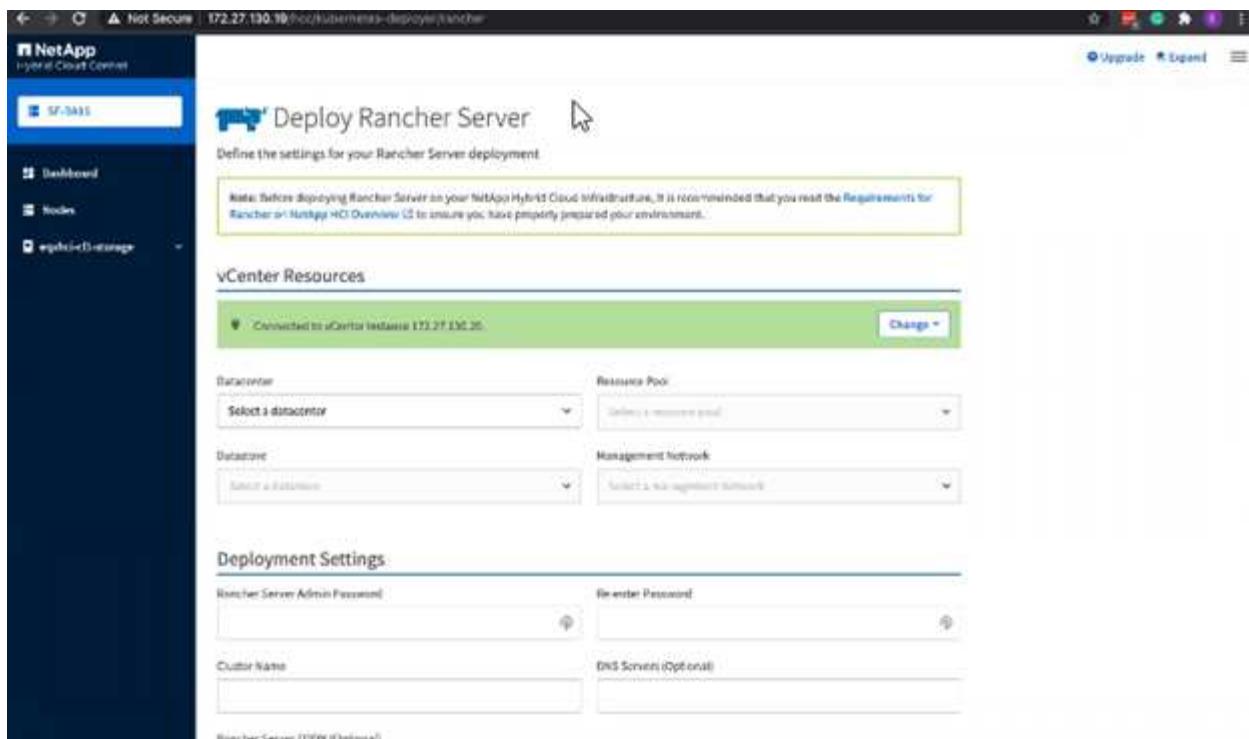


A popup window shows a message about getting started with Rancher.



2. Select **Deploy Rancher**.

The Rancher UI appears.



Your vCenter credentials are collected based on your NetApp Deployment Engine installation.

3. Enter **vCenter Resources** information. Some fields are described next.

- **Datacenter:** Select a datacenter. After you select the datacenter, all other fields are prepopulated, although you can change them.
- **Datastore:** Select a datastore on the NetApp HCI storage nodes. This datastore should be resilient and accessible to all of the VMware hosts. Do not select a local datastore that is accessible to only one of the hosts.
- **Management network:** This should be accessible from the management stations and from the virtual machine network where the user clusters will be hosted.

4. Enter **Deployment Settings** information:

- **DNS Servers:** Optional. If you use load balancing, enter the internal DNS server information.
- **Rancher Server FQDN:** To ensure that the Rancher Server remains available during node failures, provide a fully-qualified domain name (FQDN) that your DNS server can resolve to any of the IP addresses assigned to the Rancher Server cluster's nodes. This FQDN with the "https" prefix becomes the Rancher URL that you will use to access your Rancher implementation.

If no domain name is provided, wildcard DNS will be used instead and you will be able to access the Rancher Server using one of the URLs presented after the deployment completes.

5. Enter **Advanced Settings** information:

- **Assign Static IP Addresses:** If you enable static IP addressing, provide starting IP addresses for three IPv4 addresses in sequence, one for each management cluster virtual machine. Rancher on NetApp HCI deploys three management cluster virtual machines.
- **Configure Proxy Server:**

6. Review and select the checkbox for the Rancher End User License Agreement.

7. Review and select the checkbox to acknowledge information about Rancher software.

8. Select **Deploy**.

A bar indicates the deployment progress.



The Rancher deployment could take about 15 minutes.

When the deployment is complete, Rancher displays a message about the completion and provides a Rancher URL.



9. Record that Rancher URL that is displayed at the end of the deployment. You will use this URL to access the Rancher UI.

Verify your deployment by using vCenter Server

In your vSphere client, you can see the Rancher management cluster, which includes the three virtual machines.



Once you have finished deployment, do not modify the configuration of the Rancher server virtual machine cluster or remove the virtual machines. Rancher on NetApp HCI relies on the deployed RKE management cluster configuration to function normally.

What's next?

After deployment, you can do the following:

- [Complete post-deployment tasks](#)
- [Install Trident with Rancher on NetApp HCI](#)
- [Deploy user clusters and applications](#)
- [Manage Rancher on NetApp HCI](#)
- [Monitor Rancher on NetApp HCI](#)

Find more information

- [Rancher deployment troubleshooting](#)
- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp HCI Resources page](#)

Post-deployment tasks

Post-deployment tasks overview

After you deploy Rancher on NetApp HCI, you should continue with post-deployment activities.

- [Ensure Rancher Support parity](#)
- [Improve Rancher VM resiliency](#)
- [Configure monitoring](#)
- [Install Trident](#)
- [Enable Trident support for user clusters](#)

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Ensure Rancher Support parity

After you deploy Rancher on NetApp HCI, you need to ensure that the number of Rancher Support cores you purchased matches the number of CPU cores you are using for Rancher management VMs and user clusters.

If you purchased Rancher Support for only part of your NetApp HCI compute resources, you need to take action in VMware vSphere to ensure that Rancher on NetApp HCI and its managed user clusters are only running on hosts for which you have purchased Rancher Support. See the VMware vSphere documentation for information about how to help ensure this by confining compute workloads to specific hosts.

Find more information

- [vSphere HA and DRS Affinity Rules](#)
- [Create VM Anti-Affinity Rules](#)
- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Improve Rancher VM resiliency

After you deploy Rancher on NetApp HCI, your vSphere environment will contain three new nodes as virtual machines to host the Rancher environment. The Rancher web UI is available from each of these nodes. For full resiliency, each of the three virtual machines along with the corresponding virtual disks should reside on a different physical host after

events like power cycles and failovers.

To ensure that each VM and its resources remain on a different physical host, you can create VMware vSphere Distributed Resource Scheduler (DRS) anti-affinity rules. This is not automated as part of Rancher on NetApp HCI deployment.

For instructions on how to configure DRS anti-affinity rules, see the following VMware documentation resources:

[Create VM Anti-Affinity Rules](#)

[vSphere HA and DRS Affinity Rules](#)

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Enable monitoring

After you deploy Rancher on NetApp HCI, You can enable Active IQ storage monitoring (for SolidFire all-flash storage and NetApp HCI) and NetApp HCI compute monitoring (for NetApp HCI only) if you did not already do so during installation or upgrade.

For instructions on how to enable monitoring, see [Enable Active IQ and NetApp HCI monitoring](#).

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Install Trident

Learn about how to install Trident after you install Rancher on NetApp HCI. Trident is a storage orchestrator, which integrates with Docker and Kubernetes, as well as platforms built on these technologies, such as Red Hat OpenShift, Rancher, and IBM Cloud Private. The goal of Trident is to make the provisioning, connection, and consumption of storage transparent and frictionless for the applications. Trident is a fully supported open source project maintained by NetApp. Trident enables you to create, manage, and interact with persistent storage volumes in the standard Kubernetes format that you are familiar with.



For more information about Trident, see the [Trident documentation](#).

What you'll need

- You have installed Rancher on NetApp HCI.
- You have deployed your user clusters.
- You have configured your user cluster networks for Trident. See [Enable Trident support for user clusters](#) for instructions.
- You have completed the necessary prerequisite steps for work node preparation for Trident. See the [Trident documentation](#).

About this task

The Trident installer catalog is installed as part of the Rancher installation using NetApp Hybrid Cloud Control. In this task, you use the installer catalog to install and configure Trident.

As part of the Rancher installation, NetApp provides a node template. If you are not planning to use the node template that NetApp provides, and you want to provision on RHEL or CentOS, there might be additional requirements. If you change your worker node to RHEL or CentOS, there are several prerequisites that should be met. See the [Trident documentation](#).

Steps

1. From the Rancher UI, select a project for your user cluster.

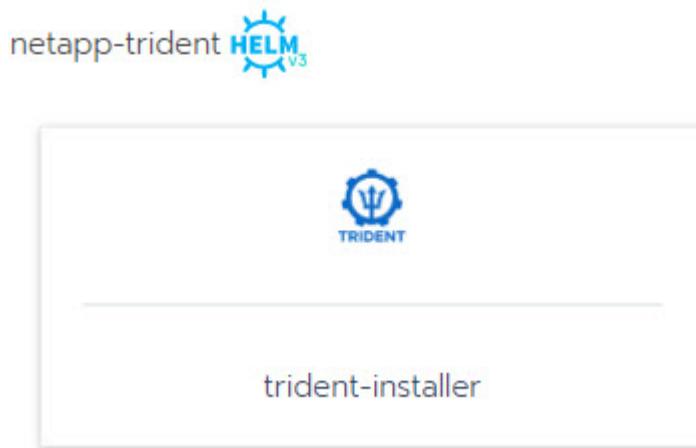


For information about projects and namespaces, see the [Rancher documentation](#).

2. Select **Apps**, and select **Launch**.



3. On the **Catalog** page, select the Trident installer.



On the page that opens, you can select the **Detailed Descriptions** arrow to learn more about the Trident app and also to find the link to the [Trident documentation](#).

4. Select the **Configurations Options** arrow, and enter the credentials and storage configuration information.

STORAGECONFIGURATION

Storage Tenant *	SVIP *
NetApp-HCI	
The name of the tenant that is already present on the SolidFire AFA.	
MVIP *	Trident Backend Name *
	solidfire
The virtual/cluster IP address for management.	
Trident Storage Driver *	
solidfire-san	The name of this Trident backend configuration.
The name of the Trident storage driver.	



The default storage tenant is NetApp HCI. You can change this value. You can also change the backend name. However, do not change the default storage driver value, which is **solidfire-san**.

5. Select **Launch**.

This installs the Trident workload on the **trident** namespace.

6. Select **Resources > Workloads**, and verify that the **trident** namespace includes the following components:

Namespace: trident		
<input type="checkbox"/>	▶	Active trident-csi
<input type="checkbox"/>	▶	Active trident-csi
<input type="checkbox"/>	▶	Active trident-installer
<input type="checkbox"/>	▶	Active trident-operator

7. (Optional) Select **Storage** for the user cluster to see the storage classes that you can use for your persistent volumes.



The three storage classes are **solidfire-gold**, **solidfire-silver**, and **solidfire-bronze**. You can make one of these storage classes the default by selecting the icon under the **Default** column.

Find more information

- [Enable Trident support for user clusters](#)
- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)

- [NetApp HCI Resources page](#)

Enable Trident support for user clusters

If your NetApp HCI environment does not have a route between the management and storage networks, and you deploy user clusters that need Trident support, you need to further configure your user cluster networks after installing Trident. For each user cluster, you need to enable communication between the management and storage networks. You can do this by modifying the networking configuration for each node in the user cluster.

About this task

Follow these general steps to modify the networking configuration for each node in the user cluster. These steps assume that you created the user cluster with the default node template that is installed with Rancher on NetApp HCI.



You can make these changes as part of a custom node template to use for future user clusters.

Steps

1. Deploy a user cluster with existing default template.
2. Connect the storage network to the user cluster.
 - a. Open the VMware vSphere web client for the connected vCenter instance.
 - b. In the Hosts and Clusters inventory tree, select a node in the newly deployed user cluster.
 - c. Edit the node's settings.
 - d. In the settings dialog, add a new network adapter.
 - e. In the **New Network** drop down list, browse for a network and select **HCI_Internal_Storage_Data_Network**.
 - f. Expand the network adapter section and record the MAC address for the new network adapter.
 - g. Select **OK**.
3. In Rancher, download the SSH private key file for each node in the user cluster.
4. Connect using SSH to a node in the user cluster, using the private key file that you have downloaded for that node:

```
ssh -i <private key filename> <ip address>
```

5. As the superuser, edit and save the `/etc/netplan/50-cloud-init.yaml` file so that it includes the `ens224` section, similar to the following example. Replace `<MAC address>` with the MAC address you recorded earlier:

```
network:
  ethernets:
    ens192:
      dhcp4: true
      match:
        macaddress: 00:50:56:91:1d:41
      set-name: ens192
    ens224:
      dhcp4: true
      match:
        macaddress: <MAC address>
      set-name: ens224
  version: 2
```

6. Use the following command to reconfigure the network:

```
`netplan try`
```

7. Repeat steps 4 through 6 for each remaining node in the user cluster.
8. When you have reconfigured the network for each node in the user cluster, you can deploy applications in the user cluster that utilize Trident.

Deploy user clusters and applications

After deploying Rancher on NetApp HCI, you can set up user clusters and add applications to those clusters.

Deploy user clusters

After deployment, Dev and Ops teams can then deploy their Kubernetes user clusters, similar to any Rancher deployment, on which they can deploy apps.

1. Access the Rancher UI using that URL provided to you at the end of the Rancher deployment.
2. Create user clusters. See Rancher documentation about [deploying workloads](#).
3. Provision user clusters in Rancher on NetApp HCI. See Rancher documentation about [setting up Kubernetes clusters in Rancher](#).

Deploy applications on user clusters

Similar to any Rancher deployment, you add applications on Kubernetes clusters.

See Rancher documentation about [deploying applications across clusters](#).

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp HCI Resources page](#)

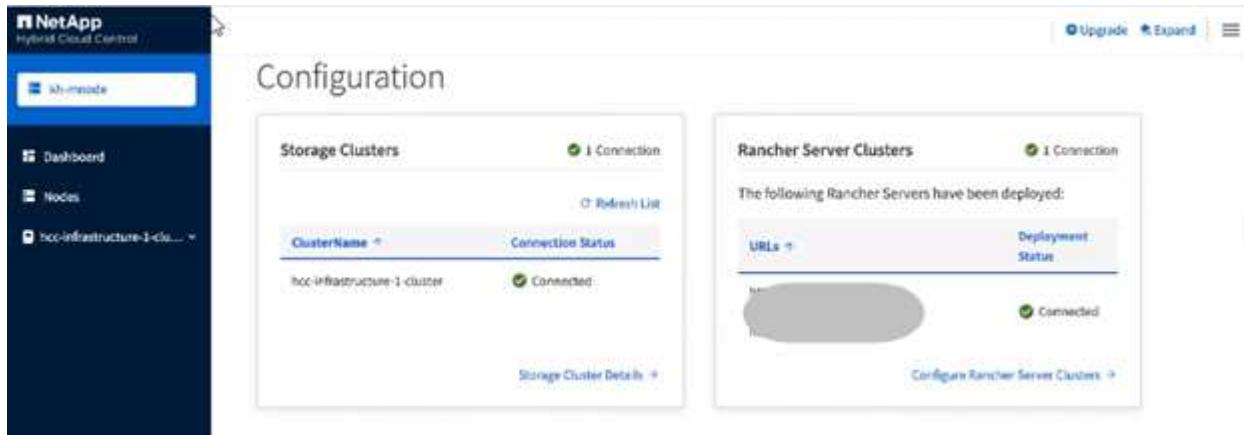
Manage Rancher on NetApp HCI

After deploying Rancher on NetApp HCI, you can view the Rancher server cluster URLs and status. You can also delete the Rancher server.

Identify Rancher server cluster URLs and status

You can identify Rancher server cluster URLs and determine server status.

1. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI or Element storage cluster administrator credentials.
2. From the Dashboard, select the top right Options icon and select **Configure**.



The screenshot shows the 'Configuration' page in the NetApp Hybrid Cloud Control interface. On the left, there is a sidebar with options like 'Dashboard', 'Nodes', and 'NetApp Infrastructure-1 Cluster'. The main area is titled 'Configuration' and contains two sections: 'Storage Clusters' and 'Rancher Server Clusters'. The 'Storage Clusters' section shows one cluster named 'NetApp Infrastructure-1 Cluster' with a 'Connected' status. The 'Rancher Server Clusters' section shows a list of deployed Rancher servers with their URLs and deployment status. One server is listed as 'Connected'.

The Rancher Server Clusters page displays a list of Rancher server clusters that have been deployed, the associated URL, and status.

Find more information

- [Remove Rancher](#)
- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp HCI Resources page](#)

Monitor a Rancher on NetApp HCI implementation

There are multiple ways to monitor Rancher server, management clusters, and other details.

- NetApp Hybrid Cloud Control
- Rancher UI

- NetApp Active IQ
- vCenter Server

Monitor Rancher using the NetApp Hybrid Cloud Control

Using the NetApp Hybrid Cloud Control, you can view the Rancher URL and Rancher server cluster status. You can also monitor the nodes in which Rancher is running.

1. Log in to NetApp Hybrid Cloud Control by providing Element storage cluster administrator credentials.
2. From the Dashboard, select the top right Options icon and select **Configure**.

3. To view nodes information, from the Hybrid Cloud Control Dashboard, expand the name of your storage cluster and select **Nodes**.

Monitor Rancher using the Rancher UI

Using the Rancher UI, you can see information about Rancher on NetApp HCI management clusters and user clusters.



In the Rancher UI, management clusters are referred to as "local clusters."

1. Access the Rancher UI using that URL provided to you at the end of the Rancher deployment.
2. See [Monitoring in Rancher v2.5](#).

Monitor Rancher using NetApp Active IQ

Using NetApp Active IQ, you can view Rancher telemetry, such as installation information, nodes, clusters, status, namespace information, and more.

1. Log in to NetApp Hybrid Cloud Control by providing Element storage cluster administrator credentials.
2. From the top right menu, select **NetApp Active IQ**.

Monitor Rancher using vCenter Server

Using vCenter Server, you can monitor the Rancher virtual machines.

Find more information

- [Rancher documentation about architecture](#)
- [Kubernetes terminology for Rancher](#)
- [NetApp Element Plug-in for vCenter Server](#)
- [NetApp HCI Resources page](#)

Upgrade Rancher on NetApp HCI

To upgrade Rancher software, you can use the NetApp Hybrid Cloud Control (HCC) UI or REST API. HCC provides an easy button process to upgrade the components of your Rancher deployment, including Rancher server, Rancher Kubernetes Engine (RKE), and the management cluster's node OS (for security updates). You can alternatively use the API to help automate upgrades.

Upgrades are available by component instead of a cumulative package. As such, some component upgrades such as the Ubuntu OS come available on a more rapid cadence. Upgrades affect only your Rancher server instance and the management cluster that Rancher Server is deployed on. Upgrades to the management cluster node's Ubuntu OS are for critical security patches only and do not upgrade the operating system. User clusters cannot be upgraded from NetApp Hybrid Cloud Control.

What you'll need

- **Admin privileges:** You have storage cluster administrator permissions to perform the upgrade.
- **Management services:** You have updated your management services bundle to the latest version.



You must upgrade to the latest management services bundle 2.17 or later for Rancher functionality.

- **System ports:** If you are using NetApp Hybrid Cloud Control for upgrades, you have ensured that the necessary ports are open. See [Network ports](#) for more information.
- **End User License Agreement (EULA):** Beginning with management services 2.20.69, you must accept and save the EULA before using the NetApp Hybrid Cloud Control UI or API to upgrade a Rancher deployment:

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. The EULA pops up. Scroll down, select **I accept for current and all future updates**, and select **Save**.

Upgrade options

Choose one of the following upgrade processes:

- [Use NetApp Hybrid Cloud Control UI to upgrade a Rancher deployment](#)
- [Use NetApp Hybrid Cloud Control API to upgrade a Rancher deployment](#)

Use NetApp Hybrid Cloud Control UI to upgrade a Rancher deployment

Using the NetApp Hybrid Cloud Control UI, you can upgrade any of these components in your Rancher deployment:

- Rancher server
- Rancher Kubernetes Engine (RKE)
- Node OS security updates

What you'll need

- A good internet connection. Dark site upgrades are not available.

Steps

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. Select **Upgrade** near the top right of the interface.
4. On the **Upgrades** page, select **Rancher**.
5. Select the **Actions** menu for the software you want to upgrade.
 - Rancher server
 - Rancher Kubernetes Engine (RKE)
 - Node OS security updates
6. Select **Upgrade** for Rancher server or RKE upgrades or **Apply Upgrade** for Node OS security updates.



For node OS, unattended upgrades for security patches are run on a daily basis but the node is not rebooted automatically. By applying upgrades, you are rebooting each node for the security updates to take effect.

A banner appears indicating the component upgrade is successful. There could be up to a 15 minute delay before NetApp Hybrid Cloud Control UI shows the updated version number.

Use NetApp Hybrid Cloud Control API to upgrade a Rancher deployment

You can use APIs to upgrade any of these components in your Rancher deployment:

- Rancher server
- Rancher Kubernetes Engine (RKE)
- Node OS (for security updates)

You can use an automation tool of your choice to run the APIs or the REST API UI available on the management node.

Options

- [Upgrade Rancher Server](#)

- [Upgrade RKE](#)
- [Apply node OS security updates](#)



For node OS, unattended upgrades for security patches are run on a daily basis but the node is not rebooted automatically. By applying upgrades, you are rebooting each node for the security updates to take effect.

Upgrade Rancher Server

API commands

1. Initiate the list upgrade versions request:

```
curl -X POST "https://<ManagementNodeIP>/k8sdeployer/1/upgrade/rancher-versions" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer \${TOKEN} used by the API command when you [authorize](#). The bearer \${TOKEN} is in the curl response.

2. Get task status using task ID from previous command and copy the latest version number from the response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

3. Initiate Rancher server upgrade request:

```
curl -X PUT "https://<mNodeIP>/k8sdeployer/1/upgrade/rancher/<version number>" -H "accept: application/json" -H "Authorization: Bearer"
```

4. Get task status using task ID from upgrade command response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

REST API UI steps

1. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/k8sdeployer/api/
```

2. Select **Authorize** and complete the following:

- a. Enter the cluster user name and password.
 - b. Enter the client ID as `mnode-client`.
 - c. Select **Authorize** to begin a session.
 - d. Close the authorization window.
3. Check for the latest upgrade package:
 - a. From the REST API UI, run **POST /upgrade/rancher-versions**.
 - b. From the response, copy the task ID.
 - c. Run **GET /task/{taskID}** with the task ID from the previous step.
 4. From the **/task/{taskID}** response, copy the latest version number you want to use for the upgrade.
 5. Run the Rancher Server upgrade:
 - a. From the REST API UI, run **PUT /upgrade/rancher/{version}** with the latest version number from the previous step.
 - b. From the response, copy the task ID.
 - c. Run **GET /task/{taskID}** with the task ID from the previous step.

The upgrade has finished successfully when the `PercentComplete` indicates 100 and `results` indicates the upgraded version number.

Upgrade RKE

API commands

1. Initiate the list upgrade versions request:

```
curl -X POST "https://<mNodeIP>/k8sdeployer/1/upgrade/rke-versions" -H
"accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer `${TOKEN}` used by the API command when you [authorize](#). The bearer `${TOKEN}` is in the curl response.

2. Get task status using task ID from previous command and copy the latest version number from the response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept:
application/json" -H "Authorization: Bearer ${TOKEN}"
```

3. Initiate the RKE upgrade request

```
curl -X PUT "https://<mNodeIP>/k8sdeployer/1/upgrade/rke/<version
number>" -H "accept: application/json" -H "Authorization: Bearer"
```

4. Get task status using task ID from upgrade command response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

REST API UI steps

1. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/k8sdeployer/api/
```

2. Select **Authorize** and complete the following:

- a. Enter the cluster user name and password.
- b. Enter the client ID as `mnode-client`.
- c. Select **Authorize** to begin a session.
- d. Close the authorization window.

3. Check for the latest upgrade package:

- a. From the REST API UI, run **POST /upgrade/rke-versions**.
- b. From the response, copy the task ID.
- c. Run **GET /task/{taskID}** with the task ID from the previous step.

4. From the **/task/{taskID}** response, copy the latest version number you want to use for the upgrade.

5. Run the RKE upgrade:

- a. From the REST API UI, run **PUT /upgrade/rke/{version}** with the latest version number from the previous step.
- b. Copy the task ID from the response.
- c. Run **GET /task/{taskID}** with the task ID from the previous step.

The upgrade has finished successfully when the `PercentComplete` indicates 100 and `results` indicates the upgraded version number.

Apply node OS security updates

API commands

1. Initiate the check upgrades request:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/upgrade/checkNodeUpdates" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```



You can find the bearer `${TOKEN}` used by the API command when you [authorize](#). The bearer `${TOKEN}` is in the curl response.

2. Get task status using task ID from previous command and verify a more recent version number is available from the response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

3. Apply the node updates:

```
curl -X POST "https://<mNodeIP>/k8sdeployer/1/upgrade/applyNodeUpdates" -H "accept: application/json" -H "Authorization: Bearer"
```



For node OS, unattended upgrades for security patches are run on a daily basis but the node is not rebooted automatically. By applying upgrades, you are rebooting each node sequentially for the security updates to take effect.

4. Get task status using task ID from the upgrade applyNodeUpdates response:

```
curl -X GET "https://<mNodeIP>/k8sdeployer/1/task/<taskID>" -H "accept: application/json" -H "Authorization: Bearer ${TOKEN}"
```

REST API UI steps

1. Open the management node REST API UI on the management node:

```
https://<ManagementNodeIP>/k8sdeployer/api/
```

2. Select **Authorize** and complete the following:

- Enter the cluster user name and password.
- Enter the client ID as `mnode-client`.
- Select **Authorize** to begin a session.
- Close the authorization window.

3. Verify if an upgrade package is available:

- From the REST API UI, run **GET /upgrade/checkNodeUpdates**.
- From the response, copy the task ID.
- Run **GET /task/{taskID}** with the task ID from the previous step.
- From the **/task/{taskID}** response, verify that there is a more recent version number than the one currently applied to your nodes.

4. Apply the node OS upgrades:



For node OS, unattended upgrades for security patches are run on a daily basis but the node is not rebooted automatically. By applying upgrades, you are rebooting each node sequentially for the security updates to take effect.

- From the REST API UI, run **POST /upgrade/applyNodeUpdates**.

- b. From the response, copy the task ID.
- c. Run **GET /task/{taskID}** with the task ID from the previous step.
- d. From the **/task/{taskID}** response, verify that the upgrade has been applied.

The upgrade has finished successfully when the `PercentComplete` indicates 100 and `results` indicates the upgraded version number.

Find more information

- [NetApp Element Plug-in for vCenter Server](#)

Remove a Rancher installation on NetApp HCI

If you accidentally deploy Rancher on NetApp HCI with incorrect information (such as an incorrect Rancher server FQDN), you need to remove the installation and then redeploy. Follow these steps to remove the Rancher installation on NetApp HCI instance.

This action does not delete the user clusters.

 You might want to retain the user clusters. If you do retain them, you can later migrate them to another Rancher implementation. If you want to delete the user clusters, you should do that first before deleting the Rancher server; otherwise, deleting the user clusters after the Rancher server is deleted is more difficult.

Options

- [Remove Rancher on NetApp HCI using NetApp Hybrid Cloud Control \(Recommended\)](#)
- [Remove Rancher on NetApp HCI using the REST API](#)

Remove Rancher on NetApp HCI using NetApp Hybrid Cloud Control

You can use NetApp Hybrid Cloud Control web UI to remove the three virtual machines that were set up during deployment to host the Rancher server.

Steps

1. Open the IP address of the management node in a web browser:

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

2. Log in to NetApp Hybrid Cloud Control by providing the storage cluster administrator credentials.
3. From the Dashboard, select the menu on the upper right.
4. Select **Configure**.

5. In the **Rancher Server Clusters** pane, select **Configure Rancher Server Clusters**.
6. Select the **Actions** menu for the Rancher installation you need to remove.



Selecting **Delete** immediately removes the Rancher on NetApp HCI management cluster.

7. Select **Delete**.

Remove Rancher on NetApp HCI using the REST API

You can use the NetApp Hybrid Cloud Control REST API to remove the three virtual machines that were set up during deployment to host the Rancher server.

Steps

1. Enter the management node IP address followed by `/k8sdeployer/api/`:

```
https://[IP address]/k8sdeployer/api/
```

2. Select **Authorize** or any lock icon and enter cluster admin credentials for permissions to use APIs.
 - a. Enter the cluster user name and password.
 - b. Select **Request body** from the type drop-down list if the value is not already selected.
 - c. Enter the client ID as `mnode-client` if the value is not already populated.
 - d. Do not enter a value for the client secret.
 - e. Select **Authorize** to begin a session.
 - f. Close the window.
3. Close the **Available authorizations** dialog box.
4. Select **POST/destroy**.
5. Select **Try it out**.
6. In the request body text box, enter the Rancher server FQDN as the `serverURL` value.
7. Select **Execute**.

After several minutes, the Rancher server virtual machines should no longer be visible in the Hosts and Clusters list in vSphere Client. After removal, you can use NetApp Hybrid Cloud Control to redeploy

Find more Information

- [Rancher deployment troubleshooting](#)
- [NetApp Element Plug-in for vCenter Server](#)

Maintain H-series hardware

H-series hardware maintenance overview

You should perform hardware maintenance tasks, such as replace faulty nodes, replace faulty drives in the storage nodes, and so on to ensure that your system functions optimally.

Here are the links to the hardware maintenance tasks:

- [Replace 2U H-series chassis](#)
- [Replace DC power supply units in H615C and H610S nodes](#)
- [Replace DIMMs in compute nodes](#)
- [Replace drives for storage nodes](#)
- [Replace H410C nodes](#)
- [Replace H410S nodes](#)
- [Replace H610C and H615C nodes](#)
- [Replace H610S nodes](#)
- [Replace power supply units](#)
- [Replace SN2010, SN2100, and SN2700 switches](#)
- [Replace storage node in a two-node cluster](#)

Find more information

- [NetApp Element Plug-in for vCenter Server](#)
- [*TR-4820: NetApp HCI Networking Quick Planning Guide*](#)
- [NetApp Configuration Advisor 5.8.1 or later network validation tool](#)

Replace 2U H-series chassis

If your chassis has a fan failure or a power issue, you should replace it as soon as possible. The steps in the chassis replacement procedure depend on your NetApp HCI configuration and cluster capacity, which requires careful consideration and planning. You should contact NetApp Support for guidance and to order a replacement chassis.

About this task

You should consider the following before you replace the chassis:

- Does your rack have additional space for a new chassis?
- Do any of the chassis in your deployment have unused node slots?
- If your rack has additional space, can you move each of the nodes from the failed chassis to the new chassis, one at a time? You should keep in mind that this process might take time.
- Can your storage cluster remain online when you remove the nodes that are part of the failed chassis?

- Can your virtual machines (VMs) and ESXi cluster handle the workload when you remove the compute nodes that are part of the failed chassis?

Replacement options

Choose from one of the following options below:

[Replace the chassis when additional unused space is available in the rack](#)

[Replace the chassis when additional unused space is not available in the rack](#)

Replace the chassis when additional unused space is available in the rack

If your rack has additional space, you can install the new chassis and move nodes one at a time to the new chassis. If any of the installed chassis have unused node slots, you can move nodes from the failed chassis to the unused slots one at a time, and then remove the failed chassis. Before you begin the procedure, ensure that the cable lengths are sufficient and switch ports are available.

 The steps for moving compute nodes are different from the steps for moving storage nodes. You should ensure that the nodes are correctly shut down before you move them. After you move all the nodes from the failed chassis, you should remove the chassis from the rack and return it to NetApp.

Install the new chassis

You can install the new chassis into the rack space available, and move the nodes into it.

What you'll need

- You have an electrostatic discharge (ESD) wristband, or you have taken other antistatic protection.
- You have the replacement chassis.
- You have a lift or two or more persons to perform the steps.
- You have a #1 Phillips screwdriver.

Steps

1. Put on antistatic protection.
2. Unpack the replacement chassis.
Keep the packaging for when you return the failed chassis to NetApp.
3. Insert the rails that were shipped to you along with the chassis.
4. Slide the replacement chassis into the rack.



Always use sufficient manpower or a lift while installing the chassis.

5. Secure the chassis to the rack with the front mounting thumb screws, and tighten the screws with the screwdriver.

Move a compute node

Before you move a compute node to the new chassis or to an existing chassis that has additional unused slots, you should migrate the virtual machines (VMs), shut the node down correctly, and label the cables inserted into the node.



Ensure that you have antistatic protection when you move the node.

Steps

1. Make a note of the serial number of the node from the sticker at the back of the node.
2. In the VMware vSphere Web Client, select **Hosts and Clusters**, select a node (host), and then select **Monitor > Hardware Status > Sensors**.
3. In the **Sensors** section, look for the serial number that you noted from the sticker at the back of the node.
4. After you find the matching serial number, migrate the VMs to another available host.



See the VMware documentation for the migration steps.

5. Right-click the node, and select **Power > Shut Down**.
You are now ready to physically removing the node from the chassis.
6. Label the node and all the cables at the back of the node.
7. Remove the node from the chassis by pulling down the cam handle on the right side of each node, and pulling the node out using both the cam handles.
8. Reinstall the node into the new chassis by pushing the node in until you hear a click.
The labels you had attached to the node before you removed it help guide you. The node powers on automatically when you install it correctly.



Ensure that you support the node from under when you install it. Do not use excessive force while pushing the node into the chassis.



If installing into the new chassis, ensure that you install the node into its original slot in the chassis.

9. Reconnect the cables to the same ports at the back of the node.
The labels you had on the cables when you disconnected them help guide you.
10. Ensure that you do not force the cables into the ports; you might damage the cables, ports, or both.
11. Confirm that the compute node (host) is listed in the ESXi cluster in the VMware vSphere Web Client.
12. Perform these steps for all the compute nodes in the failed chassis.

Move a storage node

Before you move the storage nodes to the new chassis, you should remove the drives, shut down the nodes correctly, and label all the components.

Steps

1. Identify the node that you are going to remove as follows:
 - a. Note down the serial number of the node from the sticker at the back of the node.
 - b. In the VMware vSphere Web Client, select **NetApp Element Management**, and copy the MVIP IP address.
 - c. Use the MVIP IP address in a web browser to log in to the NetApp Element software UI with the user name and password that you configured in the NetApp Deployment Engine.
 - d. Select **Cluster > Nodes**.

- e. Match the serial number you noted down with the serial number (service tag) listed.
 - f. Make a note of the node ID of the node.
2. After you identify the node, move iSCSI sessions away from the node by using the following API call:
- ```
wget --no-check-certificate -q --user=<USER> --password=<PASS> -O - --post -data '{ "method": "MovePrimariesAwayFromNode", "params": { "nodeID": <NODEID> } } 'https://<MVIP>/json-rpc/8.0
```
- MVIP is the MVIP IP address, NODEID is the node ID, USER is the user name you configured in the NetApp Deployment Engine when you set up NetApp HCI, and PASS is the password you configured in the NetApp Deployment Engine when you set up NetApp HCI.
3. Select **Cluster > Drives** to remove the drives associated with the node.
-  You should wait for the drives that you removed to show up as Available before you remove the node.
4. Select **Cluster > Nodes > Actions > Remove** to remove the node.
5. Use the following API call to shut down the node:
- ```
wget --no-check-certificate -q --user=<USER> --password=<PASS> -O - --post -data '{ "method": "Shutdown", "params": { "option": "halt", "nodes": [ <NODEID> ] } } 'https://<MVIP>/json-rpc/8.0
```
- MVIP is the MVIP IP address, NODEID is the node ID, USER is the user name you configured in the NetApp Deployment Engine when you set up NetApp HCI, and PASS is the password you configured in the NetApp Deployment Engine when you set up NetApp HCI.
- After the node is shut down, you are ready to physically remove it from the chassis.
6. Remove the drives from the node in the chassis as follows:
- a. Remove the bezel.
 - b. Label the drives.
 - c. Open the cam handle, and slide each drive out carefully using both hands.
 - d. Place the drives on an antistatic, level surface.
7. Remove the node from the chassis as follows:
- a. Label the node and cables attached to it.
 - b. Pull down the cam handle on the right side of each node, and pull the node out using both the cam handles.
8. Reinstall the node into the chassis by pushing the node in until you hear a click.
The labels you had attached to the node before you removed it help guide you.
-  Ensure that you support the node from under when you install it. Do not use excessive force while pushing the node into the chassis.
-  If installing into the new chassis, ensure that you install the node into its original slot in the chassis.
9. Install the drives into their respective slots in the node by pressing down the cam handle on each drive until it clicks.
10. Reconnect the cables to the same ports at the back of the node.
The labels you had attached to the cables when you disconnected them will help guide you.



Ensure that you do not force the cables into the ports; you might damage the cables, ports, or both.

11. After the node powers on, add the node to the cluster.



It might take up to 15 minutes for the node to get added and be displayed under **Nodes > Active**.

12. Add the drives.

13. Perform these steps for all the storage nodes in the chassis.

Replace the chassis when additional unused space is not available in the rack

If your rack does not have additional space and if none of the chassis in your deployment has unused node slots, you should determine what can stay online, if anything, before you do the replacement procedure.

About this task

You should take the following points into consideration before you do the chassis replacement:

- Can your storage cluster remain online without the storage nodes in the failed chassis?
If the answer is no, you should shut down all the nodes (both compute and storage) in your NetApp HCI deployment.
If the answer is yes, you can shut down only the storage nodes in the failed chassis.
- Can your VMs and ESXi cluster stay online without the compute nodes in the failed chassis?
If the answer is no, you must shut down or migrate the appropriate VMs to be able to shut down the compute nodes in the failed chassis.
If the answer is yes, you can shut down only the compute nodes in the failed chassis.

Shut down a compute node

Before you move the compute node to the new chassis, you should migrate the VMs, shut it down correctly, and label the cables inserted into the node.

Steps

1. Make a note of the serial number of the node from the sticker at the back of the node.
2. In the VMware vSphere Web Client, select **Hosts and Clusters**, select a node (host), and then select **Monitor > Hardware Status > Sensors**.
3. In the **Sensors** section, look for the serial number that you noted from the sticker at the back of the node.
4. After you find the matching serial number, migrate the VMs to another available host.



See the VMware documentation for the migration steps.

5. Right-click the node, and select **Power > Shut Down**.
You are now ready to physically removing the node from the chassis.

Shut down a storage node

See the steps [here](#).

Remove the node

You should ensure that you remove the node carefully from the chassis and label all the components. The steps to physically remove the node are the same for both storage and compute nodes. For a storage node, remove the drive before you remove the node.

Steps

1. For a storage node, remove the drives from the node in the chassis as follows:
 - a. Remove the bezel.
 - b. Label the drives.
 - c. Open the cam handle, and slide each drive out carefully using both hands.
 - d. Place the drives on an antistatic, level surface.
2. Remove the node from the chassis as follows:
 - a. Label the node and cables attached to it.
 - b. Pull down the cam handle on the right side of each node, and pull the node out using both the cam handles.
3. Perform these steps for all the nodes you want to remove.

You are now ready to remove the failed chassis.

Replace the chassis

If your rack does not have additional space, you should uninstall the failed chassis and replace it with the new chassis.

Steps

1. Put on antistatic protection.
2. Unpack the replacement chassis, and keep it on a level surface.
Keep the packaging for when you return the failed unit to NetApp.
3. Remove the failed chassis from the rack, and place it on a level surface.



Use sufficient manpower or a lift while moving a chassis.

4. Remove the rails.
5. Install the new rails that were shipped to you with the replacement chassis.
6. Slide the replacement chassis into the rack.
7. Secure the chassis to the rack with the front mounting thumb screws, and tighten the screws with the screwdriver.
8. Install the nodes into the new chassis as follows:
 - a. Reinstall the node into its original slot in the chassis by pushing the node in until you hear a click. The labels you attached to the node before you removed it help guide you.



Ensure that you support the node from under when you install it. Do not use excessive force while pushing the node into the chassis.

- ol style="list-style-type: none;">- b. For storage nodes, install the drives into their respective slots in the node by pressing down the cam handle on each drive until it clicks.

- c. Reconnect the cables to the same ports at the back of the node.
 The labels you attached to the cables when you disconnected them help guide you.



Ensure that you do not force the cables into the ports; you might damage the cables, ports, or both.

9. Ensure that the nodes are online as follows:

Option	Steps
If you reinstalled all the nodes (both storage and compute) in your NetApp HCI deployment	<ol style="list-style-type: none"> 1. In the VMware vSphere Web Client, confirm that the compute nodes (hosts) are listed in the ESXi cluster. 2. In the Element plug-in for vCenter server, confirm that the storage nodes are listed as Active.
If you reinstalled only the nodes in the failed chassis	<ol style="list-style-type: none"> 1. In the VMware vSphere Web Client, confirm that the compute nodes (hosts) are listed in the ESXi cluster. 2. In the Element plug-in for vCenter server, select Cluster > Nodes > Pending. 3. Select the node, and select Add. <p> It might take up to 15 minutes for the node to get added and be displayed under Nodes > Active.</p> <ol style="list-style-type: none"> 4. Select Drives. 5. From the Available list, add the drives. 6. Perform these steps for all the storage nodes you reinstalled.

10. Verify that the volumes and datastores are up and accessible.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace DC power supply units in H615C and H610S nodes

H615C and H610S nodes support two –48 V to –60 V DC power supply units. These units are available as optional add-ons when you order H615C or H610S nodes. You can use the these instructions to remove the AC power supply units in the chassis and replace them with DC power supply units, or to replace a faulty DC power supply unit with a new DC power supply unit.

What you'll need

- If you are replacing a faulty DC power supply unit, you have procured a replacement DC power supply unit.
- If you are swapping out the AC power supply units in your chassis with DC units, you have taken into consideration the downtime for the procedure.
- You have an electrostatic discharge (ESD) wristband, or you have taken other antistatic precautions.
- You have ensured that the power supply requirements are met:
 - Supply voltage: $-(48-60)$ V DC
 - Current consumption: 37A (maximum)
 - Breaker requirements: 40A breaker
- You have ensured that the materials in your environment adhere to the RoHS specifications.
- You have ensured that the cable requirements are met:
 - One UL 10 AWG, 2 m maximum (stranded) black cable $[-(48-60)$ V DC]
 - One UL 10 AWG, 2 m maximum (stranded) red cable [V DC return]
 - One UL 10 AWG, 2 m maximum green/yellow cable, green with a yellow stripe, stranded wire (safety ground)

About this task

The procedure applies to the following node models:

- One rack unit (1U) H615C compute chassis
- 1U H610S storage chassis



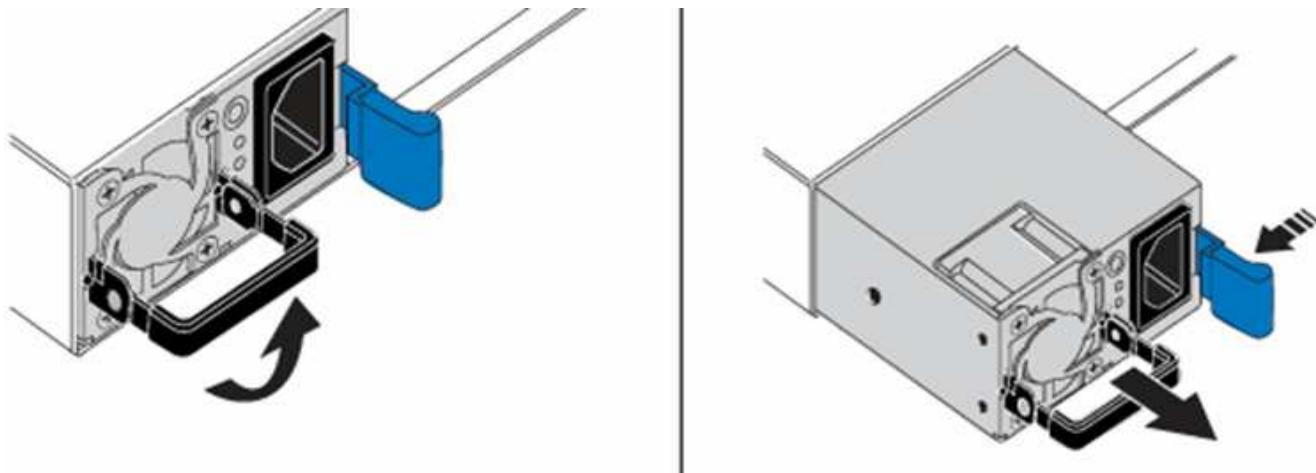
In the case of H615C and H610S, the terms "node" and "chassis" are used interchangeably because node and chassis are not separate components, unlike in the case of the 2U, four-node chassis.



You cannot mix AC and DC power supply units in your installation.

Steps

1. Turn off the power supply units and unplug the power cords. If you are replacing a faulty DC power supply unit, turn off the power source and remove all the cables inserted into the blue connector.
2. Lift the cam handle, and press the blue latch to slide out the power supply unit.



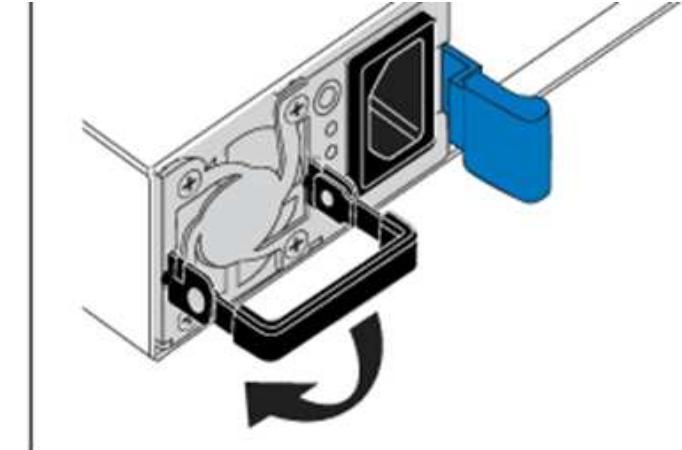
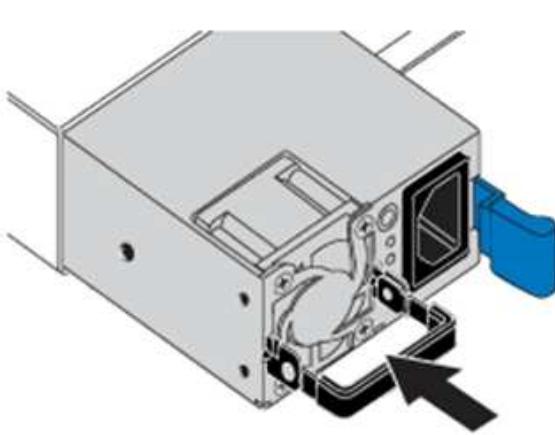


The illustration is an example. The location of the power supply unit in the chassis and the color of the release button vary depending on the type of chassis you have.



Ensure that you use both hands to support the weight of the power supply unit.

3. Using both hands, align the edges of the power supply unit with the opening in the chassis, gently push the unit into the chassis using the cam handle until it locks into place, and return the cam handle to the upright position.



4. Cable the DC power supply units. Ensure that the power source is off while cabling the DC power supply unit and the power source.
 - a. Insert the black, red, and green/yellow cables to the blue connectors.
 - b. Insert the blue connector to the DC power supply units and the power source.



5. Power on the DC power supply units.



The power supply LEDs are lit when the DC power supply unit comes online. Green LED lights indicate that the power supply units are working correctly.

6. Return the faulty unit to NetApp by following the instructions in the box that was shipped to you.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace DIMMs in compute nodes

You can replace a faulty dual inline memory module (DIMM) in NetApp HCI compute nodes instead of replacing the entire node.

What you'll need

- Before starting this procedure, you should have contacted NetApp Support and received a replacement part. Support will be involved during the installation of the replacement. If you have not done so already, contact [Support](#).
- You have planned for system downtime, because you need to power down or cycle the node and boot the node to NetApp Safe Mode to access the terminal user interface (TUI).

About this task

This procedure applies to the following compute node models:

- H410C nodes. An H410C node is inserted into a 2U NetApp HCI chassis.
- H610C node. An H610C node is built into the chassis.
- H615C node. An H615C node is built into the chassis.



H410C and H615C nodes include DIMMs from different vendors. Ensure that you do not mix DIMMs from different vendors in one chassis.



The terms "chassis" and "node" are used interchangeably in the case of H610C and H615C, because the node and chassis are not separate components.

Here are the steps involved in replacing DIMMs in compute nodes:

- [Prepare to replace the DIMM](#)
- [Replace the DIMM from the chassis](#)

Prepare to replace the DIMM

When issues with the DIMM occur, VMware ESXi displays alerts, such as Memory Configuration Error, Memory Uncorrectable ECC, Memory Transition to Critical, and Memory Critical Overtemperature. Even if the alerts disappear after a while, the hardware problem might persist. You should diagnose and address the faulty DIMM. You can get information about the faulty DIMM from vCenter Server. If you need more information than what is available from vCenter Server, you must run the hardware check in the TUI.

Steps

1. Access the node by logging in to vCenter Server.
2. Right-click the node that is reporting the error, and select the option to place the node in maintenance mode.

3. Migrate the virtual machines (VMs) to another available host.



See the VMware documentation for the migration steps.

4. Power down the compute node.



If you have the information about which DIMM needs to be replaced and do not need to access the TUI, you can skip the following steps in this section.

5. Plug in a keyboard, video, and mouse (KVM) to the back of the node that reported the error.

6. Press the power button at the front of the node.

It takes approximately six minutes for the node to boot. The screen displays a boot menu.

7. Identify the slot that logged the error as follows:

a. For H615C, do the following:

i. log in to the BMC UI.

ii. Select **Logs & Reports > IPMI Event Log**.

iii. In the event log, find the memory error and identify the slot on which the error is logged.

8. For H410C and H615C nodes, perform the steps to identify the DIMM manufacturer part number.



H410C and H615C nodes include DIMMs from different manufacturers. You should not mix different DIMM types in the same chassis. You should identify the manufacturer of the faulty DIMM and order a replacement of the same type.

a. Log in to the BMC to launch the console on the node.

b. Press **F2** on the keyboard to get to the **Customize System/View Logs** menu.

c. Enter the password when prompted.



The password should match what you configured in the NetApp Deployment Engine when you set up NetApp HCI.



- d. From the System Customization menu, press the down arrow to navigate to Troubleshooting Options, and press **Enter**.



- e. From the Troubleshooting Mode Options menu, use the up or down arrow to enable ESXi shell and SSH, which are disabled by default.
- f. Press the **<Esc>** key twice to exit Troubleshooting Options.
- g. Run the `smbiosDump` command using one of the following options:

Option	Steps
Option A	<ol style="list-style-type: none"> 1. Connect to the ESXi host (compute node) using the IP address of the host and the root credentials that you defined. 2. Run the <code>smbiosDump</code> command. See the following sample output: <pre> `Memory Device:#30 Location: "P1-DIMMA1" Bank: "P0_Node0_Channel0_Dimm0" Manufacturer:"Samsung" Serial: "38EB8380" Asset Tag: "P1-DIMMA1_AssetTag (date:18/15)" Part Number: "M393A4K40CB2-CTD" Memory Array: #29 Form Factor: 0x09 (DIMM) Type: 0x1a (DDR4) Type Detail: 0x0080 (Synchronous) Data Width: 64 bits (+8 ECC bits) Size: 32 GB`</pre>
Option B	<ol style="list-style-type: none"> 1. Press Alt + F1 to enter shell, and log in to the node to run the command.

9. Contact NetApp Support for help with the next steps. NetApp Support requires the following information to process a part replacement:

- Node serial number
- Cluster name
- System event log from the BMC UI (**Logs & Reports > IPMI Event Log> Download Event Logs**)
- Output from the `smbiosDump` command

Replace the DIMM from the chassis

Before you physically remove and replace the faulty DIMM in the chassis, ensure that you have performed all the [preparatory steps](#).



DIMMs should be replaced in the same slots they were removed from.

Steps

1. Power down the chassis or node.



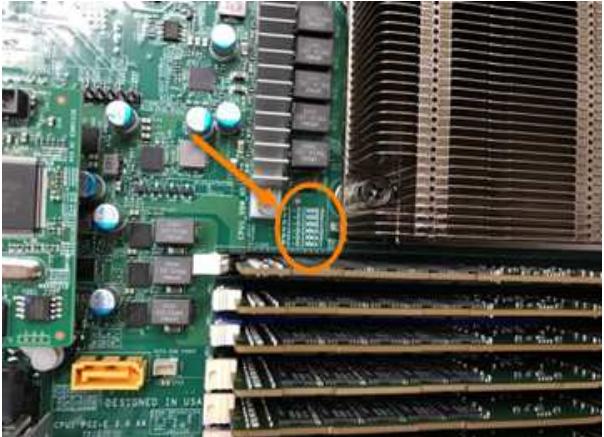
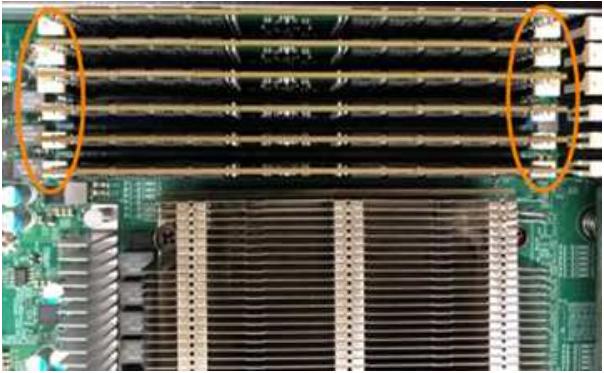
For a H610C or H615C chassis, power down the chassis. For H410C nodes in a 2U, four-node chassis, power down only the node with the faulty DIMM.

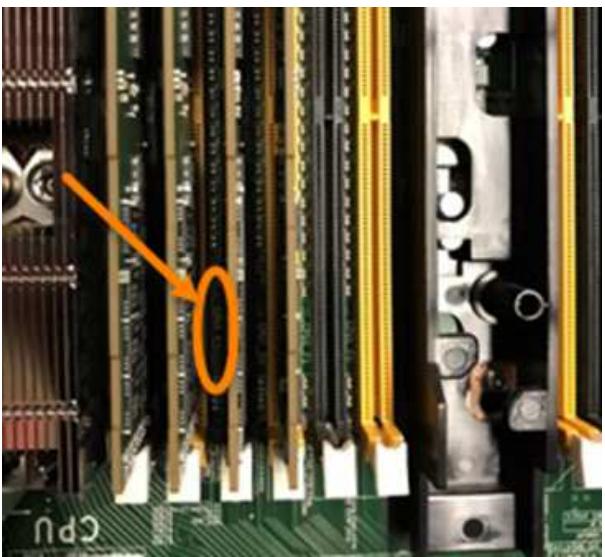
2. Remove the power cables and network cables, carefully slide the node or chassis out of the rack, and place it on a flat, antistatic surface.

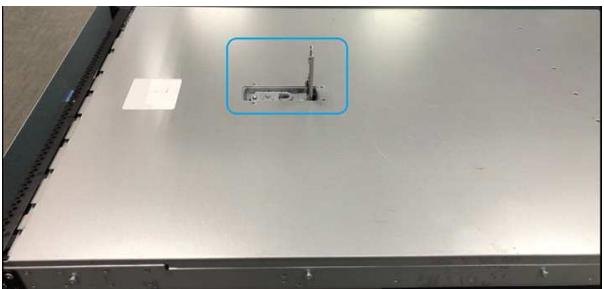
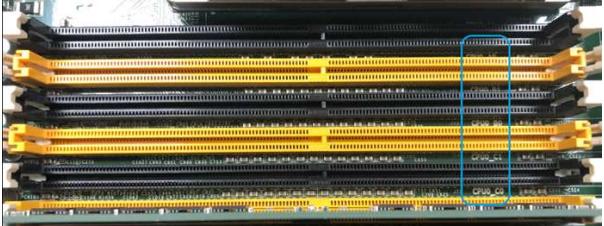


Consider using twist ties for cables.

3. Put on antistatic protection before you open the chassis cover to replace the DIMM.
4. Perform the steps relevant to your node model:

Node model	Steps
H410C	<p>1. Find the failed DIMM by matching the slot number/ID you noted earlier with the numbering on the motherboard. Here are sample images showing the DIMM slot numbers on the motherboard:</p>   <p>2. Press the two retaining clips outward, and carefully pull the DIMM up. Here is a sample image showing the retaining clips:</p>  <p>3. Install the replacement DIMM correctly. When you insert the DIMM into the slot correctly, the two clips lock in place.</p> <p> Ensure that you touch only the rear ends of the DIMM. If you press on other parts of the DIMM, it might result in damage to the hardware.</p> <p>Install the node in the NetApp HCI chassis, ensuring that the node clicks when you slide it into place.</p>

Node model	Steps
H610C	<p>1. Lift the cover as shown in the following image:</p>  <p>2. Loosen the four blue lock screws at the back of the node. Here is a sample image showing the location of two lock screws; you will find the other two on the other side of the node:</p>  <p>3. Remove both PCI card blanks.</p> <p>4. Remove the GPU and the airflow cover.</p> <p>5. Find the failed DIMM by matching the slot number/ID you noted earlier with the numbering on the motherboard. Here is a sample image showing the location of the DIMM slot numbers on the motherboard:</p>  <p>Press the two retaining clips outward, and carefully pull the DIMM up.</p>

Node model	Steps
H615C	<ol style="list-style-type: none"> 1. Lift the cover as shown in the following image:  2. Remove the GPU (if your H615C node has GPU installed) and the airflow cover.  3. Find the failed DIMM by matching the slot number/ID you noted earlier with the numbering on the motherboard. Here is a sample image showing the location of the DIMM slot numbers on the motherboard:  4. Press the two retaining clips outward, and carefully pull the DIMM up. 5. Install the replacement DIMM correctly. When you insert the DIMM into the slot correctly, the two clips lock in place. <p> Ensure that you touch only the rear ends of the DIMM. If you press on other parts of the DIMM, it might result in damage to the hardware.</p> <ol style="list-style-type: none"> 6. Replace the airflow cover. <p>Put the cover back on the node.</p> <p>Install the H610C chassis in the rack, ensuring that the chassis clicks when you slide it into place.</p>

5. Insert the power cables and network cables.
Ensure that all the port lights turn on.
6. Press the power button at the front of the node if it does not power on automatically when you install it.
7. After the node is displayed in vSphere, right-click the name and take the node out of maintenance mode.
8. Verify the hardware information as follows:
 - a. Log in to the baseboard management controller (BMC) UI.
 - b. Select **System > Hardware Information**, and check the DIMMs listed.

What's next

After the node returns to normal operation, in vCenter, check the Summary tab to ensure that the memory capacity is as expected.

-  If the DIMM is not installed correctly, the node will operate normally but with lower than expected memory capacity.
-  After the DIMM replacement procedure, you can clear the warnings and errors on the Hardware Status tab in vCenter. You can do this if you want to erase the history of errors related to the hardware that you replaced. [Learn more](#).

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace drives for storage nodes

If a drive is faulty or if the drive wear level falls below a threshold, you should replace it. Alarms in the Element software UI and VMware vSphere Web Client notify you when a drive has failed or is going to fail. You can hot-swap a failed drive.

About this task

This procedure is for replacing drives in H410S and H610S storage nodes. Removing a drive takes the drive offline. Any data on the drive is removed and migrated to other drives in 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.

Best practices for handling drives

You should follow these best practices for handling drives:

- Keep the drive in the ESD bag until you are ready to install it.
- Open the ESD bag by hand or cut the top off with a pair of scissors.
- Always wear an ESD wrist strap grounded to an unpainted surface on your chassis.
- Always use both hands when removing, installing, or carrying a drive.
- Never force a drive into the chassis.
- Always use approved packaging when shipping drives.

- Do not stack drives on top of each other.

Best practices for adding and removing drives

You should follow these best practices for adding drives to the cluster and removing drives from the cluster:

- Add all the block drives and ensure that block syncing is complete before you add the slice drives.
- For Element software 10.x and later, add all the block drives at once. Ensure that you don't do this for more than three nodes at once.
- For Element software 9.x and earlier, add three drives at once allowing them to completely sync before adding the next group of three.
- Remove the slice drive and ensure that slice syncing is complete before removing the block drives.
- Remove all the block drives from a single node at once. Ensure that all block syncing is complete before you move on to the next node.

Steps

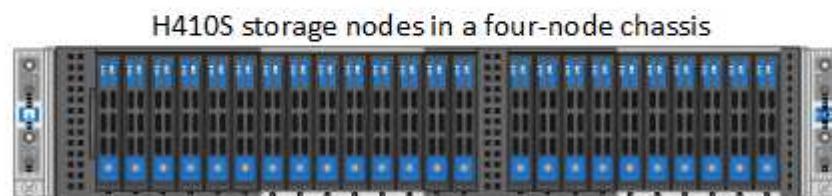
1. Remove the drive from the cluster using either the NetApp Element software UI or the NetApp Element Management extension point in Element plug-in for vCenter server.

Option	Steps
Using the Element UI	<ol style="list-style-type: none"> 1. From the Element UI, Select Cluster > Drives. 2. Select Failed to view the list of failed drives. 3. Make a note of the slot number of the failed drive. You need this information to locate the failed drive in the chassis. 4. Select Actions for the drive you want to remove. 5. Select Remove. <p>You can now physically remove the drive from the chassis.</p>

Option	Steps
Using the Element plug-in for vCenter server UI	<ol style="list-style-type: none"> 1. From the NetApp Element Management extension point of the vSphere Web Client, 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. Select the check box for each drive you want to remove. 5. Select Remove Drives. 6. Confirm the action. <p>i If there is not enough capacity to remove active drives before removing a node, an error message appears when you confirm the drive removal. After you resolve the error, you can now physically remove the drive from the chassis.</p>

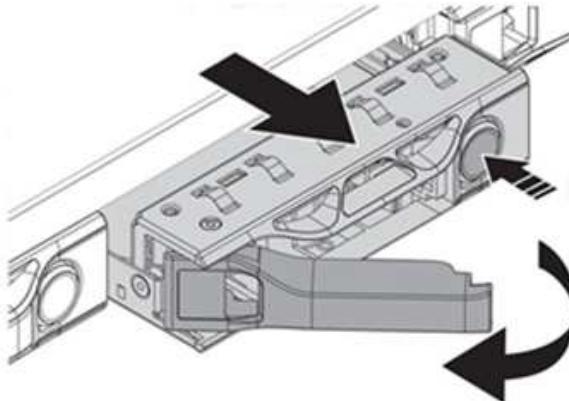
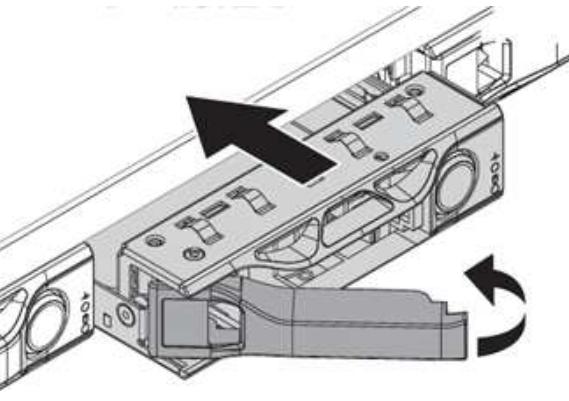
2. Replace the drive from the chassis:

- a. Unpack the replacement drive, and place it on a flat, static-free surface near the rack. Save the packing materials for when you return the failed drive to NetApp. Here is the front view of the H610S and H410S storage nodes with the drives:



- b. Perform the steps based on the node model:

Node model	Steps
H410S	<ol style="list-style-type: none"> Identify the node by matching the serial number (service tag) with the number you noted down from the Element UI. The serial number is on a sticker at the back of each node. After you identify the node, you can use the slot information to identify the slot that the failed drive is in. Drives are arranged alphabetically from A through D and from 0 through 5. Remove the bezel. Press the release button on the failed drive:  <p>When you press the release button, the cam handle on the drive springs open partially, and the drive releases from the midplane.</p> <ol style="list-style-type: none"> Open the cam handle, and slide the drive out carefully using both hands. Place the drive on an antistatic, level surface. Insert the replacement drive into the slot all the way into the chassis using both hands. Press down the cam handle until it clicks. Reinstall the bezel. Notify NetApp Support about the drive replacement. NetApp Support will provide instructions for returning the failed drive.

Node model	Steps
H610S	<ol style="list-style-type: none"> 1. Match the slot number of the failed drive from the Element UI with the number on the chassis. The LED on the failed drive is lit amber. 2. Remove the bezel. 3. Press the release button, and remove the failed drive as shown in the following illustration:  <p>i Ensure that the tray handle is fully open before you attempt to slide the drive out of the chassis.</p> <ol style="list-style-type: none"> 4. Slide the drive out, and place it on a static-free, level surface. 5. Press the release button on the replacement drive before you insert it into the drive bay. The drive tray handle springs open.  <ol style="list-style-type: none"> 6. Insert the replacement drive without using excessive force. When the drive is inserted fully, you hear a click. 7. Close the drive tray handle carefully. Reinstall the bezel. Notify NetApp Support about the drive replacement.

3. Add the drive back to the cluster using either the Element UI or the NetApp Element Management extension point in Element plug-in for vCenter server.

NetApp Support will provide instructions for returning the failed drive.



When you install a new drive in an existing node, the drive automatically registers as **Available** in the Element UI. You should add the drive to the cluster before it can participate in the cluster.

Option	Steps
Using the Element UI	<ol style="list-style-type: none">1. From the Element UI, select Cluster > Drives.2. Select Available to view the list of available drives.3. Select the Actions icon for the drive you want to add, and select Add.
Using the Element plug-in for vCenter server UI	<ol style="list-style-type: none">1. From the NetApp Element Management extension point of the vSphere Web Client, select NetApp Element Management > Cluster > Drives.2. From the Available drop-down list, select the drive, and select Add.3. Confirm the action.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace H410C nodes

You should replace a compute node in the event of CPU failure, other motherboard issues, or if it does not power on. The instructions apply to H410C nodes. If you have a H410C compute node that runs NetApp HCI Bootstrap OS version 1.6P1 or later, you do not have to replace the node if the memory DIMM fails; you need replace only the failed DIMM. If the DIMMs in your node have not failed, you can use them in the replacement node.



The replacement node should have the same version of NetApp HCI Bootstrap OS as the rest of the compute nodes in the NetApp HCI installation.

What you'll need

- You have determined that the compute node needs to be replaced.
- You have a replacement compute node.

To order a replacement node, you should contact NetApp Support. The compute node is shipped to you with the Bootstrap OS installed.

Nodes are shipped from the factory with the latest version of Bootstrap OS. You might need to perform the

return to factory image (RTFI) process on the node in the following scenarios:

- Your current NetApp HCI installation is running a version of Bootstrap OS earlier than the latest version. In this case, the RTFI process will downgrade the new node to the OS version that your NetApp HCI installation is running.
- The replacement node that is shipped is running a bootstrap OS version earlier than the latest version, and the NetApp HCI installation where the node is being replaced is already running the latest version. In this case, the RTFI process will upgrade the OS version on the new node to the latest version. See [How to RTFI using a USB key \(login required\)](#) and [How to RTFI by using the BMC \(login required\)](#).
- You have an electrostatic discharge (ESD) wristband, or you have taken other antistatic precautions.
- You have labeled each cable that is connected to the compute node.

About this task

Alarms in the VMware vSphere Web Client alert you when a node fails. You should match the serial number of the failed node from the VMware vSphere Web Client with the serial number on the sticker at the back of the node.

When replacing an H410C compute node, consider the following:

- You can intermix the H410C compute node with existing NetApp HCI compute and storage nodes in the same chassis and cluster.
- The H410C compute node operates only on high-line voltage (200-240 VAC). You should ensure that the power requirements are met when you add H410C nodes to an existing NetApp HCI system.

Steps overview

Here is a high-level overview of the steps in this procedure:

- [Prepare to replace the compute node](#)
- [Replace the compute node in the chassis](#)
- [Remove the compute node asset in NetApp HCI 1.7 and later](#)
- [Add the compute node to the cluster](#)
- [Redeploy Witness Nodes for two and three-node storage clusters](#)

Here are some additional tasks, which you might need to perform if your system has the specific conditions they are applicable to:

- [Remove Witness Nodes to free up compute resources](#)
- [Change the password if you received a replacement node with a non-standard BMC password](#)
- [Upgrade the BMC firmware on your node](#)

Prepare to replace the compute node

You should migrate the virtual machines (VMs) hosted on the node to an available host, and remove the failed node from the cluster. You should get details about the failed node, such as serial number and networking information.

Steps

1. In the VMware vSphere Web Client, perform the steps to migrate the VMs to another available host.



See the VMware documentation for the migration steps.

2. Perform the steps to remove the node from the inventory. The steps depend on the version of NetApp HCI in your current installation:

NetApp HCI version number	Steps
NetApp HCI 1.3 and later	<ol style="list-style-type: none">1. Select the failed node, and select Monitor > Hardware Status > Sensors.2. Note the serial number of the failed node. This helps you identify the node in the chassis by matching the serial number on the sticker at the back of the node with the serial number you noted.3. Right-click the failed node and select Connection > Disconnect.4. Select Yes to confirm the action.5. Right-click the failed node and select Remove from Inventory.6. Select Yes to confirm the action.
NetApp HCI versions earlier than 1.3	<ol style="list-style-type: none">1. Right-click the node and select Remove from Inventory.2. Select the failed node, and select Monitor > Hardware Status > Sensors.3. Note the Node 0 serial number, which is the serial number of the failed node. This helps you identify the node in the chassis by matching the serial number on the sticker at the back of the node with the serial number you noted.4. With the failed node selected, select Manage > Networking > VMkernel adapters, and copy the four IP addresses listed. You can reuse this information when you perform the initial network configuration steps in VMware ESXi.

Replace the compute node in the chassis

After you remove the failed node from the cluster, you can remove the node from the chassis, and install the replacement node.



Ensure that you have antistatic protection before you perform the steps here.

Steps

1. Put on antistatic protection.
2. Unpack the new node, and set it on a level surface near the chassis.

Keep the packaging material for when you return the failed node to NetApp.

3. Label each cable that is inserted at the back of the node that you want to remove.
After you install the new node, you should insert the cables back into the original ports.
4. Disconnect all the cables from the node.
5. If you want to reuse the DIMMs, remove them.
6. Pull down the cam handle on the right side of the node, and pull the node out using both the cam handles.
The cam handle that you should pull down has an arrow on it to indicate the direction in which it moves.
The other cam handle does not move and is there to help you pull the node out.



Support the node with both your hands when you pull it out of the chassis.

7. Place the node on a level surface.
You should package the node and return it to NetApp.
8. Install the replacement node.
9. Push the node in until you hear a click.



Ensure that you do not use excessive force when sliding the node into the chassis.



Ensure that the node powers on. If it does not power on automatically, push the power button at the front of the node.

10. If you removed DIMMs from the failed node earlier, insert them into the replacement node.



You should replace DIMMs in the same slots they were removed from in the failed node.

11. Reconnect the cables to the ports from which you originally disconnected them.
The labels you had attached to the cables when you disconnected them help guide you.



If the airflow vents at the rear of the chassis are blocked by cables or labels, it can lead to premature component failures due to overheating.
Do not force the cables into the ports; you might damage the cables, ports, or both.



Ensure that the replacement node is cabled in the same way as the other nodes in the chassis.

Remove the compute node asset in NetApp HCI 1.7 and later

In NetApp HCI 1.7 and later, after you physically replace the node, you should remove the compute node asset using the management node APIs. To use REST APIs, your storage cluster must be running NetApp Element software 11.5 or later and you should have deployed a management node running version 11.5 or later.

Steps

1. Enter the management node IP address followed by /mnode:
`https://[IP address]/mnode`
2. Select **Authorize** or any lock icon and enter cluster admin credentials for permissions to use APIs.
 - a. Enter the cluster user name and password.

- b. Select Request body from the type drop-down list if the value is not already selected.
- c. Enter the client ID as mnode-client if the value is not already populated.
Do not enter a value for the client secret.
- d. Select **Authorize** to begin a session.



If you get the `Auth Error TypeError: Failed to fetch` error message after you attempt to authorize, you might need to accept the SSL certificate for the MVIP of your cluster. Copy the IP in the Token URL, paste the IP into another browser tab, and authorize again. If you attempt to run a command after the token expires, you get a `Error: UNAUTHORIZED` error. If you receive this response, authorize again.

3. Close the Available authorizations dialog box.
4. Select **GET/assets**.
5. Select **Try it out**.
6. Select **Execute**.
Scroll down in the response body to the Compute section, and copy the parent and id values for the failed compute node.
7. Select **DELETE/assets/{asset_id}/compute-nodes/{compute_id}**.
8. Select **Try it out**.
Enter the parent and id values you got in step 7.
9. Select **Execute**.

Add the compute node to the cluster

You should add the compute node back to the cluster. The steps vary depending on the version of NetApp HCI you are running.

NetApp HCI 1.6P1 and later

You can use NetApp Hybrid Cloud Control only if your NetApp HCI installation runs on version 1.6P1 or later.

What you'll need

- Ensure that the vSphere instance NetApp HCI is using has vSphere Enterprise Plus licensing if you are expanding a deployment with Virtual Distributed Switches.
- Ensure that none of the vCenter or vSphere instances in use with NetApp HCI have expired licenses.
- Ensure that you have free and unused IPv4 addresses on the same network segment as existing nodes (each new node must be installed on the same network as existing nodes of its type).
- Ensure that you have the vCenter administrator account credentials ready.
- Ensure that each new node uses the same network topology and cabling as the existing storage or compute clusters.
- [Manage the initiators and volume access groups](#) for the new compute node.

Steps

1. Open the IP address of the management node in a web browser. For example:

<https://<ManagementNodeIP>>

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.
3. In the Expand Installation pane, select **Expand**.
4. Log in to the NetApp Deployment Engine by providing the local NetApp HCI storage cluster administrator credentials.



You cannot log in using Lightweight Directory Access Protocol credentials.

5. On the Welcome page, select **Yes**.
6. On the End User License page, perform the following actions:
 - a. Read the VMware End User License Agreement.
 - b. If you accept the terms, select **I accept** at the end of the agreement text.
7. Select **Continue**.
8. On the vCenter page, perform the following steps:
 - a. Enter a FQDN or IP address and administrator credentials for the vCenter instance associated with your NetApp HCI installation.
 - b. Select **Continue**.
 - c. Select an existing vSphere datacenter to which to add the new compute node, or select **Create New Datacenter** to add the new compute nodes to a new datacenter.



If you select Create New Datacenter, the Cluster field is automatically populated.

- d. If you selected an existing datacenter, select a vSphere cluster with which the new compute nodes should be associated.



If NetApp HCI cannot recognize the network settings of the cluster you have selected, ensure that the vmkernel and vmnic mapping for the management, storage, and vMotion networks are set to the deployment defaults.

- e. Select **Continue**.
9. On the ESXi Credentials page, enter an ESXi root password for the compute node or nodes you are adding.
You should use the same password that was created during the initial NetApp HCI deployment.
10. Select **Continue**.
11. If you created a new vSphere datacenter cluster, on the Network Topology page, select a network topology to match the new compute nodes you are adding.



You can only select the two-cable option if your compute nodes are using the two-cable topology and the existing NetApp HCI deployment is configured with VLAN IDs.

12. On the Available Inventory page, select the node you want to add to the existing NetApp HCI installation.



For some compute nodes, you might need to enable EVC at the highest level your vCenter version supports before you can add them to your installation. You should use the vSphere client to enable EVC for these compute nodes. After you enable it, refresh the **Inventory** page and try adding the compute nodes again.

13. Select **Continue**.
14. Optional: If you created a new vSphere datacenter cluster, on the Network Settings page, import network information from an existing NetApp HCI deployment by selecting the **Copy Setting from an Existing Cluster** checkbox.
This populates the default gateway and subnet information for each network.
15. On the Network Settings page, some of the network information has been detected from the initial deployment. The new compute node is listed by serial number, and you should assign new network information to it. For the new compute node, perform the following steps:
 - a. If NetApp HCI detected a naming prefix, copy it from the Detected Naming Prefix field, and insert it as the prefix for the new unique hostname you add in the **Hostname** field.
 - b. In the **Management IP Address** field, enter a management IP address for the compute node that is within the management network subnet.
 - c. In the vMotion IP Address field, enter a vMotion IP address for the compute node that is within the vMotion network subnet.
 - d. In the iSCSI A - IP Address field, enter an IP address for the first iSCSI port of the compute node that is within the iSCSI network subnet.
 - e. In the iSCSI B - IP Address field, enter an IP address for the second iSCSI port of the compute node that is within the iSCSI network subnet.
16. Select **Continue**.
17. On the Review page in the Network Settings section, the new node is shown in bold text. If you need to make changes to the information in any section, perform the following steps:
 - a. Select **Edit** for that section.
 - b. When finished making changes, select Continue on any subsequent pages to return to the Review page.
18. Optional: If you do not want to send cluster statistics and support information to NetApp-hosted SolidFire Active IQ servers, clear the final checkbox.
This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve problems before production is affected.
19. Select **Add Nodes**.
You can monitor the progress while NetApp HCI adds and configures the resources.
20. Optional: Verify that the new compute node is visible in vCenter.

NetApp HCI 1.4 P2, 1.4, and 1.3

If your NetApp HCI installation runs version 1.4P2, 1.4, or 1.3, you can use the NetApp Deployment Engine to add the node to the cluster.

What you'll need

- Ensure that the vSphere instance NetApp HCI is using has vSphere Enterprise Plus licensing if you are expanding a deployment with Virtual Distributed Switches.

- Ensure that none of the vCenter or vSphere instances in use with NetApp HCI have expired licenses.
- Ensure that you have free and unused IPv4 addresses on the same network segment as existing nodes (each new node must be installed on the same network as existing nodes of its type).
- Ensure that you have the vCenter administrator account credentials ready.
- Ensure that each new node uses the same network topology and cabling as the existing storage or compute clusters.

Steps

1. Browse to the management IP address of one of the existing storage nodes:

http://<storage_node_management_IP_address>/

2. Log in to the NetApp Deployment Engine by providing the local NetApp HCI storage cluster administrator credentials.



You cannot log in using Lightweight Directory Access Protocol credentials.

3. Select **Expand Your Installation**.

4. On the Welcome page, select **Yes**.

5. On the End User License page, perform the following actions:

- a. Read the VMware End User License Agreement.
- b. If you accept the terms, select **I accept** at the end of the agreement text.

6. Select **Continue**.

7. On the vCenter page, perform the following steps:

- a. Enter a FQDN or IP address and administrator credentials for the vCenter instance associated with your NetApp HCI installation.
- b. Select **Continue**.
- c. Select an existing vSphere datacenter to which to add the new compute node.
- d. Select a vSphere cluster with which the new compute node should be associated.



If you are adding a compute node with a CPU generation that is different than the CPU generation of the existing compute nodes and Enhanced vMotion Compatibility (EVC) is disabled on the controlling vCenter instance, you should enable EVC before proceeding. This ensures vMotion functionality after expansion is complete.

- e. Select **Continue**.

8. On the ESXi Credentials page, create ESXi administrator credentials for the compute node you are adding. You should use the same master credentials that were created during the initial NetApp HCI deployment.

9. Select **Continue**.

10. On the Available Inventory page, select the node you want to add to the existing NetApp HCI installation.



For some compute nodes, you might need to enable EVC at the highest level your vCenter version supports before you can add them to your installation. You should use the vSphere client to enable EVC for these compute nodes. After you enable it, refresh the Inventory page and try adding the compute nodes again.

11. Select **Continue**.
12. On the Network Settings page, perform the following steps:
 - a. Verify the information detected from the initial deployment.
 - b. Each new compute node is listed by serial number, and you should assign new network information to it. For each new storage node, perform the following steps:
 - i. If NetApp HCI detected a naming prefix, copy it from the Detected Naming Prefix field, and insert it as the prefix for the new unique hostname you add in the Hostname field.
 - ii. In the Management IP Address field, enter a management IP address for the compute node that is within the management network subnet.
 - iii. In the vMotion IP Address field, enter a vMotion IP address for the compute node that is within the vMotion network subnet.
 - iv. In the iSCSI A - IP Address field, enter an IP address for the first iSCSI port of the compute node that is within the iSCSI network subnet.
 - v. In the iSCSI B - IP Address field, enter an IP address for the second iSCSI port of the compute node that is within the iSCSI network subnet.
 - c. Select **Continue**.
13. On the Review page in the Network Settings section, the new node is shown in bold text. If you want to make changes to information in any section, perform the following steps:
 - i. Select **Edit** for that section.
 - ii. When finished making changes, select **Continue** on any subsequent pages to return to the Review page.
14. Optional: If you do not want to send cluster statistics and support information to NetApp-hosted Active IQ servers, clear the final checkbox.
This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve problems before production is affected.
15. Select **Add Nodes**.
You can monitor the progress while NetApp HCI adds and configures the resources.
16. Optional: Verify that the new compute node is visible in vCenter.

NetApp HCI 1.2, 1.1, and 1.0

After you physically replace the node, you should add it back to the VMware ESXi cluster and perform several networking configurations so that you can use all the available functionalities.



You should have a console or keyboard, video, mouse (KVM) to perform these steps.

Steps

1. Install and configure VMware ESXi version 6.0.0 as follows:
 - a. On the remote console or KVM screen, select **Power Control > Set Power Reset**.
This restarts the node.
 - b. In the Boot Menu window that opens, select **ESXi Install** by pressing the Down Arrow key.



This window stays open for only five seconds. If you do not make the selection in five seconds, you should restart the node again.

- c. Press **Enter** to start the installation process.
- d. Complete the steps in the installation wizard.



When asked to select the disk to install ESXi on, you should select the second disk drive in the list by selecting the Down Arrow key. When asked to enter a root password, you should enter the same password that you configured in the NetApp Deployment Engine when you set up NetApp HCI.

- e. After the installation is complete, press **Enter** to restart the node.



By default, the node restarts with the NetApp HCI Bootstrap OS. You should perform a one-time configuration on the node for it to use VMware ESXi.

2. Configure VMware ESXi on the node as follows:

- a. In the NetApp HCI Bootstrap OS terminal user interface (TUI) login window, enter the following information:
 - i. User name: element
 - ii. Password: catchTheFire!
- b. Press the Down Arrow key to select **OK**.
- c. Press **Enter** to log in.
- d. In the main menu, use the Down Arrow key to select **Support Tunnel > Open Support Tunnel**.
- e. In the window that is displayed, enter the port information.



You should contact NetApp Support for this information. NetApp Support logs in to the node to set the boot configuration file and complete the configuration task.

- f. Restart the node.

3. Configure the management network as follows:

- a. Log in to VMware ESXi by entering the following credentials:
 - i. User name: root
 - ii. Password: The password you set when you installed VMware ESXi.



The password should match what you configured in the NetApp Deployment Engine when you set up NetApp HCI.

- b. Select **Configure Management Network**, and press **Enter**.
- c. Select **Network Adapters**, and press **Enter**.
- d. Select **vmnic2** and **vmnic3**, and press **Enter**.
- e. Select **IPv4 Configuration**, and press the Spacebar on the keyboard to select the static configuration option.
- f. Enter the IP address, subnet mask, and default gateway information, and press **Enter**.
You can reuse the information that you copied before you removed the node. The IP address you enter here is the Management Network IP address that you copied earlier.
- g. Press **Esc** to exit the Configure Management Network section.

- h. Select **Yes** to apply the changes.
4. Add the node (host) to the cluster and configure networking so that the node is synchronized with the other nodes in the cluster as follows:
- In the VMware vSphere Web Client, select **Hosts and Clusters**.
 - Right-click the cluster that you want to add the node to, and select **Add Host**. The wizard guides you through adding the host.

When you are asked to enter the user name and password, use the following credentials:

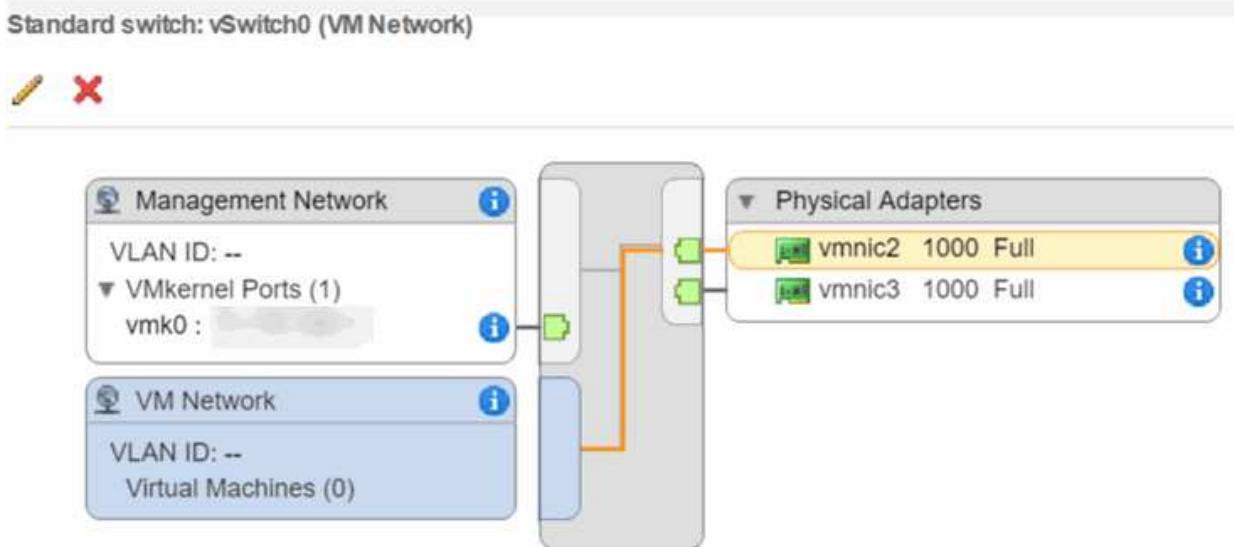


User name: root

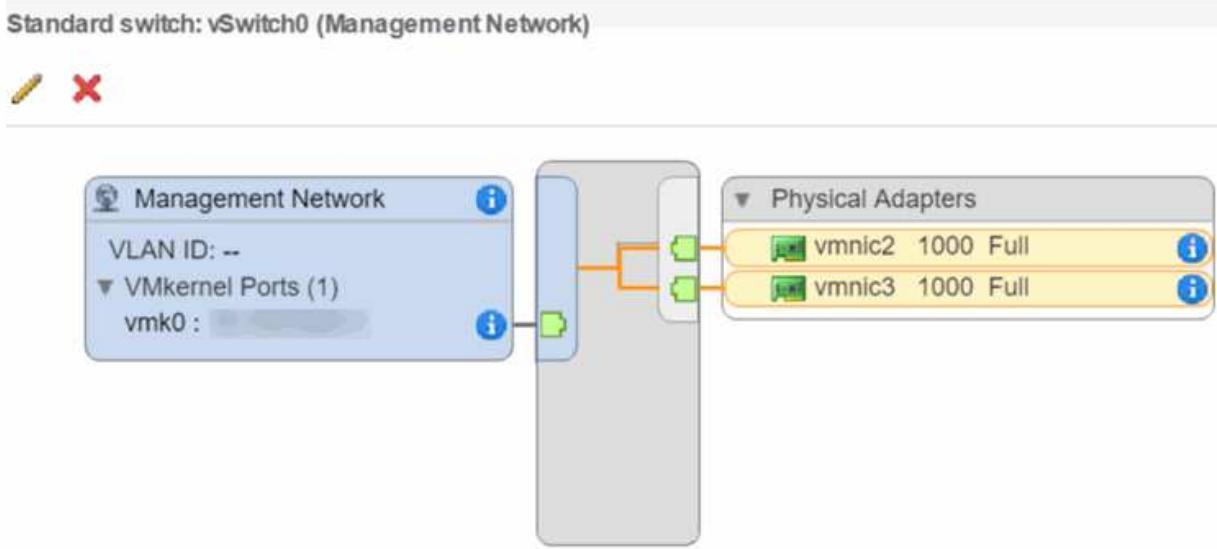
Password: The password you configured in the NetApp Deployment Engine when you set up NetApp HCI

It might take a few minutes for the node to get added to the cluster. After the process is complete, the newly added node is listed under the cluster.

- Select the node, and then select **Manage > Networking > Virtual switches**, and perform the following steps:
 - Select **vSwitch0**. You should see only vSwitch0 listed in the table that is displayed.
 - In the graphic that is displayed, select **VM Network**, and select **X** to remove the VM Network port group.



- Confirm the action.
- Select **vSwitch0**, and then select the pencil icon to edit the settings.
- In the vSwitch0 - Edit settings window, select **Teaming and failover**.
- Ensure that **vmnic3** is listed under **Standby adapters**, and select **OK**.
- In the graphic that is displayed, select **Management Network**, and select the pencil icon to edit the settings.



- viii. In the Management Network - Edit settings window, select **Teaming and failover**.
- ix. Move vmnic3 to Standby adapters by using the arrow icon, and select **OK**.
- d. From the Actions drop-down menu, select **Add Networking**, and enter the following details in the window that is displayed:
 - i. For connection type, select **Virtual Machine Port Group for a Standard Switch**, and select **Next**.
 - ii. For target device, select the option to add a new standard switch, and select **Next**.
 - iii. Select **+**.
 - iv. In the Add Physical Adapters to Switch window, select vmnic0 and vmnic4, and select **OK**. vmnic0 and vmnic4 are now listed under Active adapters.
 - v. Select **Next**.
 - vi. Under connection settings, verify that VM Network is the network label, and select **Next**.
 - vii. If you are ready to proceed, select **Finish**. vSwitch1 is displayed in the list of virtual switches.
- e. Select **vSwitch1**, and select the pencil icon to edit the settings as follows:
 - i. Under Properties, set MTU to 9000, and select **OK**. In the graphic that is displayed, select **VM Network**, and select the pencil icon to edit the settings as follows:
 - f. Select **Security**, and make the following selections:

Promiscuous mode:	<input checked="" type="checkbox"/> Override	Accept	▼
MAC address changes:	<input checked="" type="checkbox"/> Override	Reject	▼
Forged transmits:	<input checked="" type="checkbox"/> Override	Accept	▼

- i. Select **Teaming and failover**, and select the **Override** checkbox.
- ii. Move vmnic0 to Standby adapters by using the arrow icon.

- iii. Select **OK**.
- g. With vSwitch1 selected, from the Actions drop-down menu, select **Add Networking**, and enter the following details in the window that is displayed:
- For connection type, select **VMkernel Network Adapter**, and select **Next**.
 - For target device, select the option to use an existing standard switch, browse to vSwitch1, and select **Next**.
 - Under port properties, change the network label to vMotion, select the checkbox for vMotion traffic under Enable services, and select **Next**.
 - Under IPv4 settings, provide the IPv4 information, and select **Next**.
The IP address you enter here is the vMotion IP address that you copied earlier.
 - If you are ready to proceed, select **Finish**.
- h. In the graphic that is displayed, select vMotion, and select the pencil icon to edit the settings as follows:
- Select **Security**, and make the following selections:
-
- | | | | |
|----------------------|--|--------|---|
| Promiscuous mode: | <input checked="" type="checkbox"/> Override | Accept | ▼ |
| MAC address changes: | <input checked="" type="checkbox"/> Override | Reject | ▼ |
| Forged transmits: | <input checked="" type="checkbox"/> Override | Accept | ▼ |
- Select **Teaming and failover**, and select the **Override** checkbox.
 - Move vmnic4 to Standby adapters by using the arrow icon.
 - Select **OK**.
- i. With vSwitch1 selected, from the Actions drop-down menu, select **Add Networking** and enter the following details in the window that is displayed:
- For connection type, select **VMkernel Network Adapter**, and select **Next**.
 - For target device, select the option to add a new standard switch, and select **Next**.
 - Select **+**.
 - In the Add Physical Adapters to Switch window, select vmnic1 and vmnic5, and select **OK**.
vmnic1 and vmnic5 are now listed under Active adapters.
 - Select **Next**.
 - Under port properties, change the network label to iSCSI-B, and select **Next**.
 - Under IPv4 settings, provide the IPv4 information, and select **Next**.
The IP address you enter here is the iSCSI-B IP address that you copied earlier.
 - If you are ready to proceed, select **Finish**.
vSwitch2 is displayed in the list of virtual switches.
- j. Select **vSwitch2**, and select the pencil icon to edit the settings as follows:
- Under Properties, set MTU to 9000, and select **OK**.
- k. In the graphic that is displayed, select **iSCSI-B**, and select the pencil icon to edit the settings as follows:

- i. Select **Security**, and make the following selections:
-

Promiscuous mode:	<input checked="" type="checkbox"/> Override	Accept	▼
MAC address changes:	<input checked="" type="checkbox"/> Override	Reject	▼
Forged transmits:	<input checked="" type="checkbox"/> Override	Accept	▼

- ii. Select **Teaming and failover**, and select the **Override** checkbox.
- iii. Move vmnic1 to Unused adapters by using the arrow icon.
- iv. Select **OK**.
- I. From the Actions drop-down menu, select **Add Networking** and enter the following details in the window that is displayed:
- For connection type, select **VMkernel Network Adapter**, and select **Next**.
 - For target device, select the option to use an existing standard switch, browse to vSwitch2, and select **Next**.
 - Under port properties, change the network label to iSCSI-A, and select **Next**.
 - Under IPv4 settings, provide the IPv4 information, and select **Next**.
The IP address you enter here is the iSCSI-A IP address that you copied earlier.
 - If you are ready to proceed, select **Finish**.
- m. In the graphic that is displayed, select **iSCSI-A**, and select the pencil icon to edit the settings as follows:
- Select **Security**, and make the following selections:
-

Promiscuous mode:	<input checked="" type="checkbox"/> Override	Accept	▼
MAC address changes:	<input checked="" type="checkbox"/> Override	Reject	▼
Forged transmits:	<input checked="" type="checkbox"/> Override	Accept	▼

- Select **Teaming and failover**, and select the **Override** checkbox.
 - Move vmnic5 to Unused adapters by using the arrow icon.
 - Select **OK**.
- n. With the newly added node selected and the Manage tab open, select **Storage > Storage Adapters**, and perform the following steps:
- Select **+** and select **Software iSCSI Adapter**.
 - To add the iSCSI adapter, select **OK** in the dialog box.
 - Under Storage Adapters, select the iSCSI adapter, and from the Properties tab, copy the iSCSI Name.

Properties Devices Paths Targets Network Port Binding Advanced Options

Status Enabled

General

Name	vmhba40
Model	iSCSI Software Adapter
iSCSI Name	[REDACTED]
iSCSI Alias	



You need the iSCSI Name when you create the initiator.

- o. Perform the following steps in the NetApp SolidFire vCenter Plug-in:
 - i. Select **Management > Initiators > Create**.
 - ii. Select **Create a Single Initiator**.
 - iii. Enter the IQN address you copied earlier in the IQN/WWPN field.
 - iv. Select **OK**.
 - v. Select **Bulk Actions**, and select **Add to Volume Access Group**.
 - vi. Select **NetApp HCI**, and select **Add**.
- p. In the VMware vSphere Web Client, under Storage Adapters, select the iSCSI adapter, and perform the following steps:
 - i. Under Adapter Details, select **Targets > Dynamic Discovery > Add**.
 - ii. Enter the SVIP IP address in the iSCSI Server field.



To get the SVIP IP address, select **NetApp Element Management**, and copy the SVIP IP address.
Leave the default port number as is. It should be 3260.

- iii. Select **OK**.
A message recommending a rescan of the storage adapter is displayed.
- iv. Select the rescan icon.



- v. Under Adapter Details, select **Network Port Binding**, and select **+**.
- vi. Select the check boxes for iSCSI-B and iSCSI-A, and select **OK**.
A message recommending a rescan of the storage adapter is displayed.
- vii. Select the rescan icon.
After the rescan is complete, verify if the volumes in the cluster are visible on the new compute node (host).

Redeploy Witness Nodes for two and three-node storage clusters

After you physically replace the failed compute node, you should redeploy the NetApp HCI Witness Node VM if the failed compute node was hosting the Witness Node. These instructions apply only to compute nodes that are part of a NetApp HCI installation with two or three-node storage clusters.

What you'll need

- Gather the following information:
 - Cluster name from the storage cluster
 - Subnet mask, gateway IP address, DNS server, and domain information for the management network
 - Subnet mask for the storage network
- Ensure that you have access to the storage cluster to be able to add the Witness Nodes to the cluster.
- Consider the following conditions to help you decide whether to remove the existing Witness Node from VMware vSphere Web Client or the storage cluster:
 - If you want to use the same VM name for the new Witness Node, you should delete all the references to the old Witness Node from vSphere.
 - If you want to use the same host name on the new Witness Node, you should first remove the old Witness Node from the storage cluster.



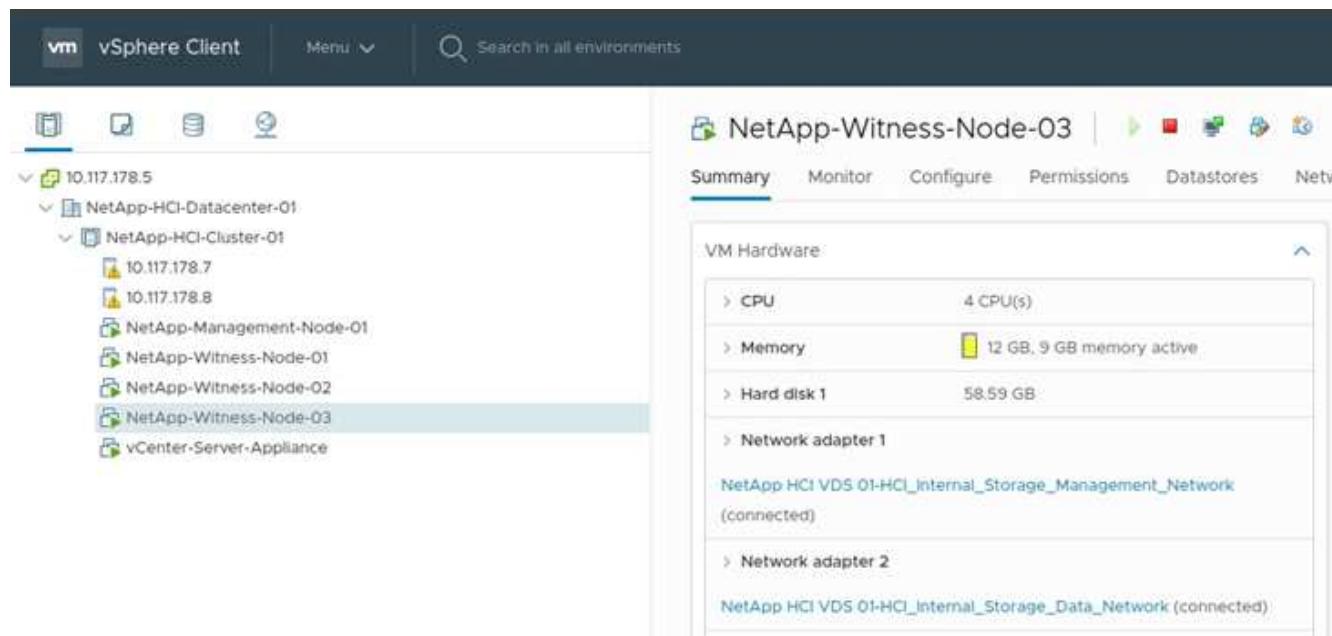
You cannot remove the old Witness Node if your cluster is down to only two physical storage nodes (and no Witness Nodes). In this scenario, you should add the new Witness Node to the cluster first before removing the old one.

You can remove the Witness Node from the cluster using the NetApp Element Management extension point.

When should you redeploy Witness Nodes?

You should redeploy Witness Nodes in the following scenarios:

- You replaced a failed compute node that is part of a NetApp HCI installation, which has a two or three-node storage cluster and the failed compute node was hosting a Witness Node VM.
- You performed the return to factory image (RTFI) procedure on the compute node.
- The Witness Node VM is corrupted.
- The Witness Node VM was accidentally removed from ESXi.
The VM is configured using the template that is created as part of initial deployment using the NetApp Deployment Engine. Here is an example of what a Witness Node VM looks like:



Steps

1. In the VMware vSphere Web Client, select **Hosts and Clusters**.
2. Right-click the compute node that will host the Witness Node VM, and select **New Virtual Machine**.
3. Select **Deploy from template**, and select **Next**.
4. Follow the steps in the wizard:
 - a. Select **Data Center**, locate the VM template, and select **Next**.
 - b. Enter a name for the VM in the following format: NetApp-Witness-Node-##
5. Leave the default selection for VM location as is, and select **Next**.
6. Leave the default selection for the destination compute resource as is, and select **Next**.
7. Select the local datastore, and select **Next**.
Free space on the local datastore varies depending on the compute platform.
8. Select **Power on virtual machine after creation** from the list of deploy options, and select **Next**.
9. Review the selections, and select **Finish**.
10. Configure the management and storage network, and cluster settings for the Witness Node as follows:
 - a. In the VMware vSphere Web Client, select **Hosts and Clusters**.
 - b. Right-click the Witness Node, and power it on if it is not already powered on.
 - c. In the Summary view of the Witness Node, select **Launch Web Console**.
 - d. Wait for the Witness Node to boot up to the menu with the blue background.
 - e. Select anywhere inside the console to access the menu.
 - f. Configure the management network as follows:
 - i. Press the down arrow key to navigate to Network, and then press **Enter** for OK.
 - ii. Navigate to **Network config**, and then press **Enter** for OK.

- iii. Navigate to **net0**, and then press **Enter** for OK.
- iv. Press **Tab** till you get to the IPv4 field, and then if applicable, delete the existing IP in the field and enter the management IP information for the Witness Node. Check the subnet mask and gateway as well.



No VLAN tagging will be applied at the VM host level; tagging will be handled in vSwitch.

- v. Press **Tab** to navigate to OK, and press **Enter** to save changes.
After management network configuration, the screen returns to Network.

g. Configure the storage network as follows:

- i. Press the down arrow key to navigate to Network, and then press **Enter** for OK.
- ii. Navigate to **Network config**, and then press **Enter** for OK.
- iii. Navigate to **net1**, and then press **Enter** for OK.
- iv. Press **Tab** till you get to the IPv4 field, and then if applicable, delete the existing IP in the field and enter the storage IP information for the Witness Node.
- v. Press **Tab** to navigate to OK, and press **Enter** to save the changes.
- vi. Set MTU to 9000.



If MTU is not set before you add the Witness Node to the cluster, you see cluster warnings for inconsistent MTU settings. This can prevent garbage collection from running and cause performance problems.

- vii. Press **Tab** to navigate to OK, and press **Enter** to save changes.
After storage network configuration, the screen returns to Network.

h. Configure the cluster settings as follows:

- i. Press **Tab** to navigate to Cancel, and press **Enter**.
- ii. Navigate to **Cluster settings**, and then press **Enter** for OK.
- iii. Press **Tab** to navigate to Change Settings, and press **Enter** for Change Settings.
- iv. Press **Tab** to navigate to Hostname field, and enter the host name.
- v. Press the down arrow key to access the Cluster field and enter the cluster name from the storage cluster.
- vi. Press the **Tab** key to navigate to OK button, and press **Enter**.

6. Add the Witness Node to the storage cluster as follows:

- a. From the vSphere Web Client, access the NetApp Element Management extension point from the **Shortcuts** tab or the side panel.
- b. Select **NetApp Element Management > Cluster**.
- c. Select the **Nodes** sub-tab.
- d. Select **Pending** from the drop-down list to view the list of nodes.
The Witness Node should appear in the pending nodes list.
- e. Select the check box for the node you want to add, and select **Add node**.
When the action is complete, the node appears in the list of active nodes for the cluster.

Change the password if you received a replacement node with a non-standard BMC password

Some replacement nodes may be shipped with non-standard passwords for the baseboard management controller (BMC) UI. If you receive a replacement node with a non-standard BMC password, you should change the password to the default, ADMIN.

Steps

1. Identify whether you received a replacement node with a non-standard BMC password:
 - a. Look for a sticker under the IPMI port at the back of the replacement node that you received. If you locate a sticker under the IPMI port, it means that you received a node with a non-standard BMC password. See the following sample image:



- b. Make a note of the password.
2. Log in to the BMC UI using the unique password found on the sticker.
3. Select **Factory Default**, and select the **Remove current settings and set the user defaults to ADMIN/ADMIN** radio button:
4. Select **Restore**.
5. Log out and then log back in to confirm that the credentials are now changed.

Upgrade the BMC firmware on your node

After you replace the compute node, you might have to upgrade the firmware version. You can download the latest firmware file from the drop-down menu on the [NetApp Support Site \(login required\)](#).

Steps

1. Log in to the baseboard management controller (BMC) UI.
2. Select **Maintenance > Firmware Update**.

System	Server Health	Configuration	Remote Control	Virtual Media	Maintenance	Miscellaneous	Help
<ul style="list-style-type: none">✖ System⌚ FRU Reading⌚ Hardware Information	<ul style="list-style-type: none">✖ System	<p>Firmware Revision : 03.25 Firmware Build Time : 06/12/2017 BIOS Version : NA2.1 BIOS Build Time : 07/10/2017 Redfish Version : 1.0.1 CPLD Version : 01.a1.06</p>	<p>IP address : 010.063.104.248 BMC MAC address : 0:c4:7a:29:c1:d0 System LAN1 MAC address : 0:c4:7a:f3 System LAN2 MAC address : 0:c4:7a:f3 System LAN3 MAC address : 0:c4:7a:d0 System LAN4 MAC address : 0:c4:7a:d5:67:eb</p>	<ul style="list-style-type: none">Firmware UpdateUnit ResetiKVM ResetFactory DefaultIPMI ConfigurationSystem Event LogBIOS UpdateSystem Crash Dump			

3. From within the BMC console, select **Maintenance**.

Host Identification
Server: 010.116.129.202
User: root (Administrator)

System Server Health Configuration Remote Control Virtual Media Maintenance

Maintenance

Firmware Update

Unit Reset

IKVM Reset

Factory Default

IPMI Configuration

Maintenance Event Log

BIOS Update

Maintenance

Use these pages to maintain the IPMI device, such as to update firmware or reset the IPMI device.

- Firmware Update : Perform firmware update.
- Unit Reset : Reboot IPMI Device.
- IKVM Reset : Reset virtual media, IKVM keyboard & mouse.
- Factory Default : Reset IPMI to Factory defaults settings. IPMI connection will be reset.
- IPMI Configuration : Save or reload IPMI configuration.
- Maintenance Event Log : You can turn on/off the maintenance event log on this page.
- BIOS Update : Perform BIOS update.

4. From within the Maintenance tab, select **Firmware Update** from the navigation on the left of the UI, and select **Enter Update Mode**.

Host Identification
Server: 010.116.129.202
User: root (Administrator)

System Server Health Configuration Remote Control Virtual Media Maintenance

Maintenance

Firmware Update

Unit Reset

IKVM Reset

Factory Default

IPMI Configuration

Maintenance Event Log

BIOS Update

Firmware Update

This page is used to update firmware. Please note that once update mode has been initiated, the device will reset even if the update process is canceled.

Enter Update Mode

5. Select **Yes** in the confirmation dialog box.
6. Select **Browse** to select the firmware image to upload, and select **Upload Firmware**.
Loading firmware from a location outside of the direct vicinity of the node might cause extended load times and possible timeouts.
7. Allow the preserve configuration checks, and select **Start Upgrade**.
The upgrade should take approximately 5 minutes. If your upload time exceeds 60 minutes, cancel the upload and transfer the file to a local machine within the vicinity of the node.
If your session times out, you might see a number of alerts while attempting to log back in to the firmware update area of the BMC UI. If you cancel the upgrade, you are redirected to the login page.
8. After the update is complete, select **OK**, and wait for the node to reboot.
Log in after the upgrade, and select **System** to verify that the **Firmware Revision** version matches the version you uploaded.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace H410S nodes

You should replace a storage node in the event of dual inline memory module (DIMM) failure, CPU failure, Radian card problems, other motherboard issues, or if it does not power on. Alarms in the VMware vSphere Web Client alert you when a storage node is faulty. You should use the NetApp Element software UI to get the serial number (service tag) of the failed node. You need this information to locate the failed node in the chassis.

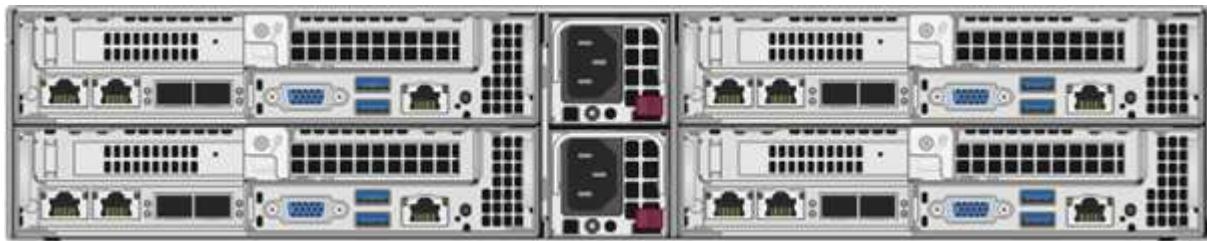
What you'll need

- You have determined that the storage node needs to be replaced.
- You have a replacement storage node.
- You have an electrostatic discharge (ESD) wristband, or you have taken other antistatic precautions.
- You have labeled each cable that is connected to the storage node.

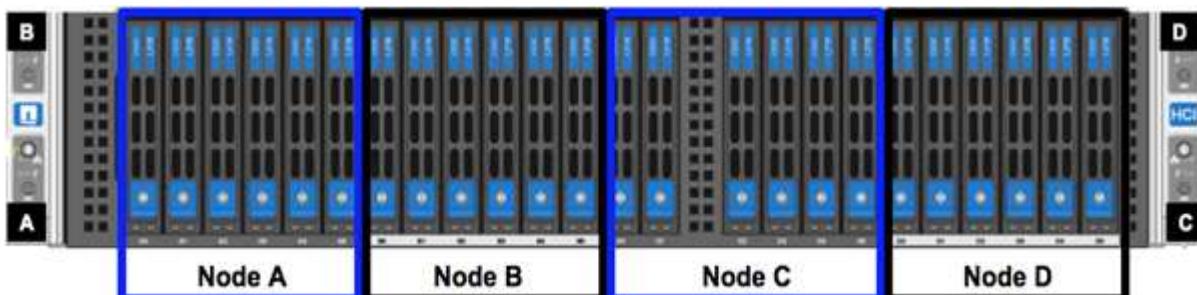
About this task

The replacement procedure applies to H410S storage nodes in a two rack unit (2U), four-node NetApp HCI chassis.

Here is the rear view of a four-node chassis with H410S nodes:



Here is the front view of a four-node chassis with H410S nodes, showing the bays that correspond to each node:



Steps overview

Here is a high-level overview of the steps in this procedure:

- [Prepare to replace the storage node](#)

- Replace the storage node in the chassis
- Add the storage node to the cluster

Prepare to replace the storage node

You should remove the faulty storage node correctly from the cluster before you install the replacement node. You can do this without causing any service interruption. You should obtain the serial number of the failed storage node from the Element UI and match it with the serial number on the sticker at the back of the node.



In the case of component failures where the node is still online and functioning, for example, a dual inline memory module (DIMM) failure, you should remove the drives from the cluster before you remove the failed node.

Steps

1. If you have a DIMM failure, remove the drives associated with the node you are going to replace from the cluster. You can use either the NetApp Element software UI or the NetApp Element Management extension point in Element plug-in for vCenter server before you remove the node.
2. Remove the nodes using either the NetApp Element software UI or the NetApp Element Management extension point in Element plug-in for vCenter server:

Option	Steps
Using the Element UI	<ol style="list-style-type: none"> 1. From the Element UI, select Cluster > Nodes. 2. Note the serial number (service tag) of the faulty node. You need this information to match it with the serial number on the sticker at the back of the node. 3. After you note the serial number, remove the node from the cluster as follows: 4. Select Actions for the node you want to remove. 5. Select Remove. <p>You can now physically remove the node from the chassis.</p>
Using the Element plug-in for vCenter server UI	<ol style="list-style-type: none"> 1. From the NetApp Element Management extension point of the vSphere Web Client, select NetApp Element Management > Cluster. 2. Select the Nodes sub-tab. 3. From Active view, select the check box for each node you want to remove, select Actions > Remove. 4. Confirm the action. Any nodes removed from a cluster appear in the list of Pending nodes.

Replace the storage node in the chassis

You should install the replacement node in the same slot in the chassis from which you remove the faulty node. You should use the serial number you noted down from the UI and match it with the serial number at the back of the node.



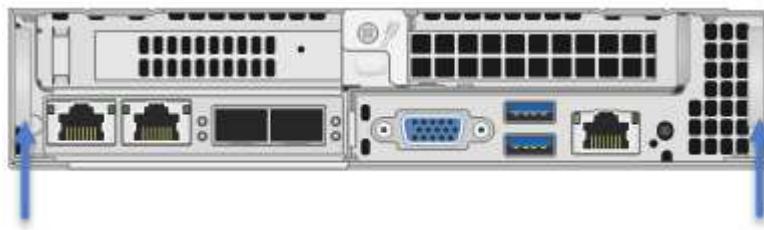
Ensure that you have antistatic protection before you perform the steps here.

Steps

1. Unpack the new storage node, and set it on a level surface near the chassis.
Keep the packaging material for when you return the failed node to NetApp.
2. Label each cable that is inserted at the back of the storage node that you want to remove.
After you install the new storage node, you must insert the cables into the original ports.
3. Disconnect all the cables from the storage node.
4. Pull down the cam handle on the right side of the node, and pull the node out using both the cam handles.
The cam handle that you should pull down has an arrow on it to indicate the direction in which it moves.
The other cam handle does not move and is there to help you pull the node out.



Support the node with both your hands when you pull it out of the chassis.



5. Place the node on a level surface.
6. Install the replacement node.
7. Push the node in until you hear a click.



Ensure that you do not use excessive force when sliding the node into the chassis.

8. Reconnect the cables to the ports from which you originally disconnected them.
The labels you had attached to the cables when you disconnected them help guide you.



If the airflow vents at the rear of the chassis are blocked by cables or labels, it can lead to premature component failures due to overheating.
Do not force the cables into the ports; you might damage the cables, ports, or both.



Ensure that the replacement node is cabled in the same way as the other nodes in the chassis.

9. Press the button at the front of the node to power it on.

Add the storage node to the cluster

You should add the storage node back to the cluster. The steps vary depending on the version of NetApp HCI you are running.

What you'll need

- You have free and unused IPv4 addresses on the same network segment as existing nodes (each new node must be installed on the same network as existing nodes of its type).
- You have one of the following types of SolidFire storage cluster accounts:
 - The native Administrator account that was created during initial deployment
 - A custom user account with Cluster Admin, Drives, Volumes, and Nodes permissions
- You have cabled and powered on the new node.
- You have the management IPv4 address of an already installed storage node. You can find the IP address in the **NetApp Element Management > Cluster > Nodes** tab of the NetApp Element Plug-in for vCenter Server.
- You have ensured that the new node uses the same network topology and cabling as the existing storage clusters.



Ensure that storage capacity is split evenly across all chassis for the best reliability.

NetApp HCI 1.6P1 and later

You can use NetApp Hybrid Cloud Control only if your NetApp HCI installation runs on version 1.6P1 or later.

Steps

1. Open the IP address of the management node in a web browser. For example:

```
https://<ManagementNodeIP>/manager/login
```

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.
3. In the Expand Installation pane, select **Expand**.
4. Log in to the NetApp Deployment Engine by providing the local NetApp HCI storage cluster administrator credentials.



You cannot log in using Lightweight Directory Access Protocol credentials.

5. On the Welcome page, select **No**.
6. Select **Continue**.
7. On the Available Inventory page, select the storage node you want to add to the existing NetApp HCI installation.
8. Select **Continue**.
9. On the Network Settings page, some of the network information has been detected from the initial deployment. Each new storage node is listed by serial number, and you should assign new network information to it. Perform the following steps:

- a. If NetApp HCI detected a naming prefix, copy it from the Detected Naming Prefix field, and insert it as the prefix for the new unique hostname you add in the Hostname field.
- b. In the Management IP Address field, enter a management IP address for the new storage node that is within the management network subnet.
- c. In the Storage (iSCSI) IP Address field, enter an iSCSI IP address for the new storage node that is within the iSCSI network subnet.
- d. Select **Continue**.



NetApp HCI might take some time to validate the IP addresses you enter. The **Continue** button becomes available when IP address validation is complete.

10. On the Review page in the Network Settings section, new nodes are shown in bold text. If you need to make changes to information in any section, perform the following steps:
 - a. Select **Edit** for that section.
 - b. When finished making changes, select **Continue** on any subsequent pages to return to the Review page.
11. Optional: If you do not want to send cluster statistics and support information to NetApp-hosted Active IQ servers, clear the final checkbox.
This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve problems before production is affected.
12. Select **Add Nodes**.
You can monitor the progress while NetApp HCI adds and configures the resources.
13. Optional: Verify that any new storage nodes are visible in the VMware vSphere Web Client.

NetApp HCI 1.4 P2, 1.4, and 1.3

If your NetApp HCI installation runs version 1.4P2, 1.4, or 1.3, you can use the NetApp Deployment Engine to add the node to the cluster.

Steps

1. Browse to the management IP address of one of the existing storage nodes:
http://<storage_node_management_IP_address>/
2. Log in to the NetApp Deployment Engine by providing the local NetApp HCI storage cluster administrator credentials.



You cannot log in using Lightweight Directory Access Protocol credentials.

3. Select **Expand Your Installation**.
4. On the Welcome page, select **No**.
5. Select **Continue**.
6. On the Available Inventory page, select the storage node to add to the NetApp HCI installation.
7. Select **Continue**.
8. On the Network Settings page, perform the following steps:
 - a. Verify the information detected from the initial deployment.
Each new storage node is listed by serial number, and you should assign new network information to it.

For each new storage node, perform the following steps:

- i. If NetApp HCI detected a naming prefix, copy it from the Detected Naming Prefix field, and insert it as the prefix for the new unique hostname you add in the Hostname field.
 - ii. In the Management IP Address field, enter a management IP address for the new storage node that is within the management network subnet.
 - iii. In the Storage (iSCSI) IP Address field, enter an iSCSI IP address for the new storage node that is within the iSCSI network subnet.
- b. Select **Continue**.
- c. On the Review page in the Network Settings section, the new node is shown in bold text. If you want to make changes to information in any section, perform the following steps:
- i. Select **Edit** for that section.
 - ii. When finished making changes, select **Continue** on any subsequent pages to return to the Review page.
9. Optional: If you do not want to send cluster statistics and support information to NetApp-hosted Active IQ servers, clear the final checkbox.
This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve problems before production is affected.
10. Select **Add Nodes**.
You can monitor the progress while NetApp HCI adds and configures the resources.
11. Optional: Verify that any new storage nodes are visible in the VMware vSphere Web Client.

NetApp HCI 1.2, 1.1, and 1.0

When you install the node, the terminal user interface (TUI) displays the fields necessary to configure the node. You must enter the necessary configuration information for the node before you proceed with adding the node to the cluster.



You must use the TUI to configure static network information as well as cluster information. If you were using out-of-band management, you must configure it on the new node.

You should have a console or keyboard, video, mouse (KVM) to perform these steps, and have the network and cluster information necessary to configure the node.

Steps

1. Attach a keyboard and monitor to the node.
The TUI appears on the tty1 terminal with the Network Settings tab.
2. Use the on-screen navigation to configure the Bond1G and Bond10G network settings for the node. You should enter the following information for Bond1G:
 - IP address. You can reuse the Management IP address from the failed node.
 - Subnet mask. If you do not know, your network administrator can provide this information.
 - Gateway address. If you do not know, your network administrator can provide this information.
You should enter the following information for Bond10G:
 - IP address. You can reuse the Storage IP address from the failed node.
 - Subnet mask. If you do not know, your network administrator can provide this information.

3. Enter **s** to save the settings, and then enter **y** to accept the changes.
4. Enter **c** to navigate to the Cluster tab.
5. Use the on-screen navigation to set the hostname and cluster for the node.



If you want to change the default hostname to the name of the node you removed, you should do it now.



It is best to use the same name for the new node as the node you replaced to avoid confusion in the future.

6. Enter **s** to save the settings.
The cluster membership changes from Available to Pending.
7. In NetApp Element Plug-in for vCenter Server, select **NetApp Element Management > Cluster > Nodes**.
8. Select **Pending** from the drop-down list to view the list of available nodes.
9. Select the node you want to add, and select **Add**.



It might take up to 15 minutes for the node to be added to the cluster and displayed under **Nodes > Active**.



Adding the drives all at once can lead to disruptions. For best practices related to adding and removing drives, see [this KB article](#) (login required).

10. Select **Drives**.
11. Select **Available** from the drop-down list to view the available drives.
12. Select the drives you want to add, and select **Add**.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace H610C and H615C nodes

You should replace a chassis to repair compute node failures related to the CPU, the motherboard, or if it does not power on. If you have a faulty DIMM in your H610C compute node that runs NetApp HCI Bootstrap OS version 1.6 or later, you can replace the DIMM and do not have to replace the chassis. For H615C nodes, you need not replace the chassis if a DIMM fails; you can replace only the failed DIMM.



For H610C and H615C, the terms "node" and "chassis" are used interchangeably, because the node and chassis are not separate components.

What you'll need

- You have verified that the node has failed.
- You have a replacement chassis.

To order a replacement, you should contact NetApp Support.

- You have an electrostatic discharge (ESD) wristband, or you have taken other antistatic protection.
- You have labeled each cable that is connected to the chassis.

About this task

Alarms in the VMware vSphere Web Client alert you when a host fails. You must match the serial number of the failed host from the VMware vSphere Web Client with the serial number on the sticker at the back of the node.

Steps overview

Here is a high-level overview of the steps in this procedure:

[Prepare to replace the node](#)

[Replace the node](#)

[Add the node to the cluster](#)

[Install the GPU drivers](#)

Prepare to replace the node

Before you replace the node, you should migrate the virtual machines (VMs) hosted on the node to an available host, and remove the node from the cluster. You should get details about the node, such as serial number and networking information.

 In the case of component failures where the node is still online and functioning, for example, a dual inline memory module (DIMM) failure, you should remove the drives from the cluster before you remove the failed node.

Steps

1. In the VMware vSphere Web Client, perform the steps to migrate the VMs to another available host.



See the VMware documentation for the migration steps.

2. Select the failed node, and select **Monitor > Hardware Status > Sensors**.
3. Make a note of the serial number of the failed node. The following screenshot is only an example:

The screenshot shows the 'Monitor' tab selected in the top navigation bar. Below it, the 'Tasks & Events' tab is also selected. The main content area is titled 'CIM data' and shows three sections: 'Sensors', 'Alerts and warnings', and 'System event log'. The 'Sensors' section displays the following details: BIOS Manufacturer: American Megatrends Inc., Model: H610C, and Serial Number: 371839000. The 'Serial Number' is highlighted with an orange circle. The 'System event log' section states 'No alerts or warnings out of 79 sensors'.

You need the serial number to identify the chassis by matching the number that you noted with the serial number on the sticker at the back of the node.

4. Right-click the failed node and select **Connection > Disconnect**.
5. Select **Yes** to confirm the action.
6. Right-click the failed node and select **Remove from Inventory**.
7. Select **Yes** to confirm the action.

Replace the node

After you remove the failed node from the cluster, you can remove the failed chassis, and install the replacement chassis.



Ensure that you have antistatic protection before you perform the steps here.

Steps

1. Unpack the new chassis, and set it on a level surface.
Keep the packaging material for when you return the failed chassis to NetApp.
2. Label each cable that is inserted at the back of the chassis that you are going to remove.
After you install the new chassis, you must insert the cables back into the original ports.
3. Disconnect all the cables from the back of the chassis.
4. Remove the chassis by unscrewing the thumbscrews on the mounting ears.
You must package and return the failed chassis to NetApp.
5. Slide the replacement chassis on to the rails.



Ensure that you do not use excessive force when sliding the chassis on to the rails.

6. Only for H615C. Remove the DIMMs from the failed chassis and insert these DIMMs in the replacement chassis.



You should replace the DIMMs in the same slots they were removed from in the failed node.

7. Remove the two power supply units on either side of the failed chassis and insert them in the replacement chassis.
8. Reconnect the cables to the ports from which you originally disconnected them.
The labels you had added on the cables when you disconnected them will help guide you.



If the airflow vents at the rear of the chassis are blocked by cables or labels, it can lead to premature component failures due to overheating.
Do not force the cables into the ports; you might damage the cables, ports, or both.

9. Power on the chassis.

Add the node to the cluster

You should configure NetApp HCI to use the new compute node.

What you'll need

- The vSphere instance NetApp HCI is using has vSphere Enterprise Plus licensing if you are adding the node to a deployment with Virtual Distributed Switches.
- None of the vCenter or vSphere instances in use with NetApp HCI have expired licenses.
- You have free and unused IPv4 addresses on the same network segment as existing nodes (the new node must be installed on the same network as existing nodes of its type).
- You have the vCenter administrator account credentials ready.

Steps

1. Open the IP address of the management node in a web browser. For example:

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

2. Log in to NetApp Hybrid Cloud Control by providing the NetApp HCI storage cluster administrator credentials.
3. In the Expand Installation pane, select **Expand**.
The browser opens the NetApp Deployment Engine.
4. Log in to the NetApp Deployment Engine by providing the local NetApp HCI storage cluster administrator credentials.



You cannot log in using Lightweight Directory Access Protocol credentials.

5. On the Welcome page, select **Yes**.
6. On the End User License page, perform the following actions:
 - a. Read the VMware End User License Agreement.
 - b. If you accept the terms, select **I accept** at the end of the agreement text.
7. Select Continue.
8. On the vCenter page, perform the following steps:

- a. Enter a FQDN or IP address and administrator credentials for the vCenter instance associated with your NetApp HCI installation.
- b. Select **Continue**.
- c. Select an existing vSphere datacenter to which to add the new compute nodes, or select **Create New Datacenter** to add the new compute nodes to a new datacenter.



If you select **Create New Datacenter**, the **Cluster** field is automatically populated.

- d. If you selected an existing datacenter, select a vSphere cluster with which the new compute nodes should be associated.



If the NetApp HCI cannot recognize the network settings of the cluster you have selected for expansion, ensure that the `vmkernel` and `vmnic` mapping for the management, storage and vMotion networks are set to the deployment defaults.

- e. Select **Continue**.
9. On the ESXi Credentials page, enter an ESXi root password for the compute node or nodes you are adding.

You should use the same password that was created during the initial NetApp HCI deployment.

10. Select **Continue**.

11. If you created a new vSphere datacenter cluster, on the Network Topology page, select a network topology to match the new compute nodes you are adding.



You can only select the two-cable option if your compute nodes are using the two-cable topology and the existing NetApp HCI deployment is configured with VLAN IDs.

12. On the Available Inventory page, select the node to add to the existing NetApp HCI installation.



For some compute nodes, you might need to enable EVC at the highest level your vCenter version supports before you can add them to your installation. You should use the vSphere client to enable EVC for these compute nodes. After you enable it, refresh the Inventory page and try adding the compute nodes again.

13. Select **Continue**.

14. Optional: If you created a new vSphere datacenter cluster, on the Network Settings page, import network information from an existing NetApp HCI deployment by selecting the **Copy Setting from an Existing Cluster** checkbox.

This populates the default gateway and subnet information for each network.

15. On the Network Settings page, some of the network information has been detected from the initial deployment. Each new compute node is listed by serial number, and you should assign new network information to it. For each new compute node, perform the following steps:

- a. If NetApp HCI detected a naming prefix, copy it from the **Detected Naming Prefix** field, and insert it as the prefix for the new unique hostname you add in the **Hostname** field.
- b. In the **Management IP Address** field, enter a management IP address for the compute node that is within the management network subnet.
- c. In the **vMotion IP Address** field, enter a vMotion IP address for the compute node that is within the vMotion network subnet.

- d. In the iSCSI A - IP Address field, enter an IP address for the first iSCSI port of the compute node that is within the iSCSI network subnet.
- e. In the iSCSI B - IP Address field, enter an IP address for the second iSCSI port of the compute node that is within the iSCSI network subnet.

16. Select **Continue**.

17. On the Review page in the Network Settings section, the new node is shown in bold text. If you need to make changes to information in any section, perform the following steps:

- a. Select **Edit** for that section.
- b. When finished making changes, select **Continue** on any subsequent pages to return to the Review page.

18. Optional: If you do not want to send cluster statistics and support information to NetApp-hosted SolidFire Active IQ servers, clear the final checkbox.

This disables real-time health and diagnostic monitoring for NetApp HCI. Disabling this feature removes the ability for NetApp to proactively support and monitor NetApp HCI to detect and resolve problems before production is affected.

19. Select **Add Nodes**.

You can monitor the progress while NetApp HCI adds and configures the resources.

20. Optional: Verify that any new compute nodes are visible in vCenter.

Install the GPU drivers

Compute nodes with NVIDIA graphics processing units (GPUs), like the H610C node, need the NVIDIA software drivers installed in VMware ESXi so that they can take advantage of the increased processing power. To install the GPU drivers, the compute node must have a GPU card.

Steps

1. Open a browser and browse to the NVIDIA licensing portal at the following URL:
<https://nvid.nvidia.com/dashboard/>
2. Download one of the following driver packages to your computer, depending on your environment:

vSphere version	Driver package
vSphere 6.0	NVIDIA-GRID-vSphere-6.0-390.94-390.96-392.05.zip
vSphere 6.5	NVIDIA-GRID-vSphere-6.5-410.92-410.91-412.16.zip
vSphere 6.7	NVIDIA-GRID-vSphere-6.7-410.92-410.91-412.16.zip

3. Extract the driver package on your computer.
The resulting .VIB file is the uncompressed driver file.
4. Copy the .VIB driver file from your computer to ESXi running on the compute node. The following example commands for each version assume that the driver is located in the \$HOME/NVIDIA/ESX6.x/ directory on the management host. The SCP utility is readily available in most Linux distributions, or available as a downloadable utility for all versions of Windows:

Option	Description
ESXi 6.0	scp \$HOME/NVIDIA/ESX6.0/NVIDIA**.vib root@<ESXi_IP_ADDR>:/.
ESXi 6.5	scp \$HOME/NVIDIA/ESX6.5/NVIDIA**.vib root@<ESXi_IP_ADDR>:/.
ESXi 6.7	scp \$HOME/NVIDIA/ESX6.7/NVIDIA**.vib root@<ESXi_IP_ADDR>:/.

5. Use the following steps to log in as root to the ESXi host and install the NVIDIA vGPU manager in ESXi.

a. Run the following command to log in to the ESXi host as the root user:

```
ssh root@<ESXi_IP_ADDRESS>
```

b. Run the following command to verify that no NVIDIA GPU drivers are currently installed:

```
nvidia-smi
```

This command should return the message nvidia-smi: not found.

c. Run the following commands to enable maintenance mode on the host and install the NVIDIA vGPU Manager from the VIB file:

```
esxcli system maintenanceMode set --enable true  
esxcli software vib install -v /NVIDIA**.vib
```

You should see the message Operation finished successfully.

d. Run the following command and verify that all eight GPU drivers are listed in the command output:

```
nvidia-smi
```

e. Run the following command to verify that the NVIDIA vGPU package was installed and loaded correctly:

```
vmkload_mod -l | grep nvidia
```

The command should return output similar to the following: nvidia 816 13808

f. Run the following commands to exit maintenance mode and reboot the host:

```
esxcli system maintenanceMode set -enable false  
reboot -f
```

6. Repeat steps 4-6 for any other newly deployed compute nodes with NVIDIA GPUs.

7. Perform the following tasks using the instructions in the NVIDIA documentation site:

a. Install the NVIDIA license server.

b. Configure the virtual machine guests for NVIDIA vGPU software.

c. If you are using vGPU-enabled desktops in a virtual desktop infrastructure (VDI) context, configure VMware Horizon View for NVIDIA vGPU software.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace H610S nodes

You might need to replace the chassis if the fan, central processing unit (CPU), or dual inline memory module (DIMM) fails, or to fix overheating issues or problems with the boot process. The blinking amber LED in the front of the chassis is an indication of a possible need for chassis replacement. You should contact NetApp Support before you proceed.

 See the [KB article](#) for information about installation requirements for H610S nodes. New and spare H610S storage nodes might have additional installation requirements based on the existing Element software version of the storage cluster. Contact NetApp Support for more information.

 The terms "node" and "chassis" are used interchangeably in the case of H610S, which is a one rack unit (1U) chassis.

Best practices for adding and removing drives

You should follow these best practices for adding drives to the cluster:

- Add all the block drives and ensure that block syncing is complete before you add the slice drives.
- For Element software 10.x and later, add all the block drives at once. Ensure that you don't do this for more than three nodes at once.
- For Element software 9.x and earlier, add three drives at once allowing them to completely sync before adding the next group of three.
- Remove the slice drive and ensure that slice syncing is complete before removing the block drives.
- Remove all the block drives from a single node at once. Ensure that all block syncing is complete before you move on to the next node.

What you'll need

- You have contacted NetApp Support.
If you are ordering a replacement, you should have a case open with NetApp Support.
- You have obtained the replacement node.
- You have an electrostatic discharge (ESD) wristband, or you have taken other antistatic protection.
- If you need to perform the Return to Factory Image (RTFI) process, you have obtained the USB key. NetApp Support can help you decide if you need to perform the RTFI process.
- You have a keyboard and monitor.
- You have removed the failed node correctly from the cluster.
- If a DIMM has failed, you have removed drives before you remove the node from the cluster.

About this task

Alarms in the VMware vSphere Web Client alert you when a host fails. You must match the serial number of the failed host from the VMware vSphere Web Client with the serial number on the sticker at the back of the node.

Steps

1. Locate the service tag at the front of the failed chassis.



2. Verify that the serial number on the service tag matches the NetApp Support case number when you ordered the replacement chassis.
3. Plug in the keyboard and monitor to the back of the failed chassis.
4. Verify the serial number of the failed node with NetApp Support.
5. Power down the chassis.
6. Label the drives in the front and cables at the back with their locations, so that you can put them back in the same locations after the replacement.

See the following image for the placement of the drives in the chassis:



7. Remove the cables.
8. Remove the chassis by unscrewing the thumbscrews on the mounting ears.
You should package and return the failed chassis to NetApp.
9. Install the replacement chassis.
10. Remove the drives carefully from the failed chassis, and insert them in the replacement chassis.



You should insert the drives in the same slots they were in before you removed them.

11. Remove the power supply units from the failed chassis, and insert them in the replacement chassis.
12. Insert the power supply cables, and the network cables in their original ports.
13. Small form-factor pluggable (SFP) transceivers might be inserted in the 10GbE ports of the replacement node. You should remove them before you cable the 10GbE ports.



See your switch vendor's documentation if your switch does not recognize the cables.

14. Power on the chassis by pressing the power button at the front.
It takes approximately five minutes and 30 seconds for the node to boot.
15. Perform the configuration steps.
 - If the H610S node is part of a NetApp HCI installation, use NetApp Hybrid Cloud Control to configure the storage resource. See [Expand NetApp HCI storage resources](#).
 - If the H610S node is part of a SolidFire all-flash storage installation, configure the node using the NetApp Element software user interface (UI).
Contact NetApp Support for assistance.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace power supply units

Each chassis includes two power supply units for power redundancy. If a power supply unit is faulty, you should replace it as soon as possible to ensure that the chassis has a redundant power source.

What you'll need

- You have determined that the power supply unit is faulty.
- You have a replacement power supply unit.
- You have verified that the second power supply unit is operating.
- You have an electrostatic discharge (ESD) wristband, or you have taken other antistatic precautions.

About this task

The replacement procedure applies to the following node models:

- Two rack unit (2U), four-node NetApp HCI chassis
- 2U H610C compute chassis
- One rack unit (1U) H615C compute chassis
- 1U H610S storage chassis



In the case of H610C, H615C, and H610S, the terms "node" and "chassis" are used interchangeably because node and chassis are not separate components, unlike in the case of the 2U, four-node chassis.

Alarms in the VMware vSphere Web Client provide information about the failed power supply unit, referring to it as PS1 or PS2. In a NetApp HCI 2U, four-node chassis, PS1 refers to the unit on the top row of the chassis and PS2 refers to the unit on the bottom row of the chassis. You can replace the faulty power supply unit while your chassis is powered on and working, as long as the redundant power supply unit is functioning.

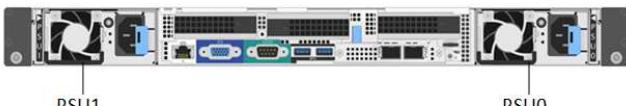
Steps

1. Locate the faulty power supply unit in the chassis. The LED on the faulty unit displays amber.

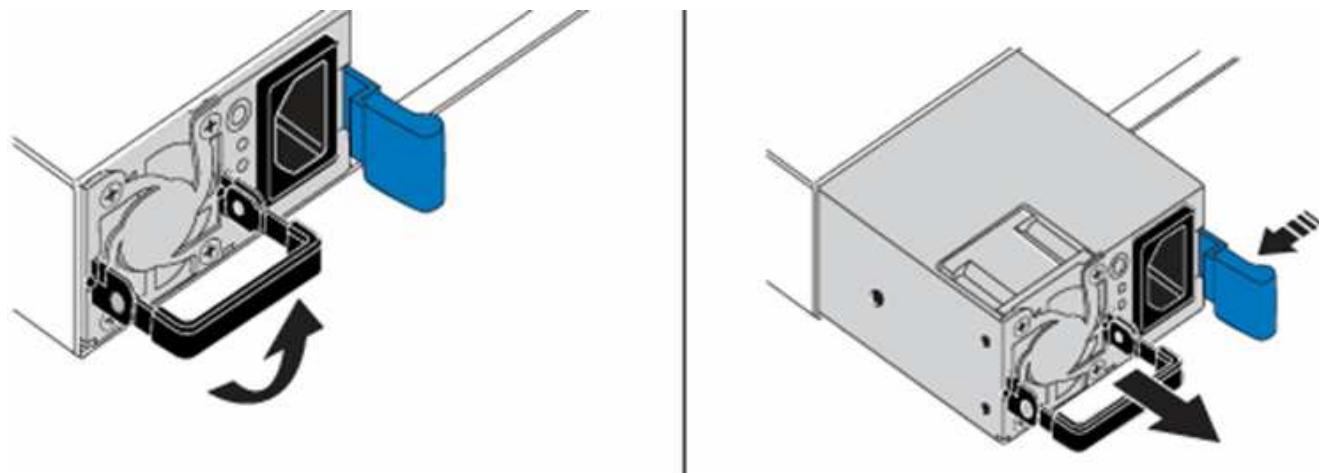


The power supply units are located differently based on the type of chassis.

See the images below for the locations of the power supply units:

Model	Location of the power supply units
2U, four-node NetApp HCI storage chassis	<p></p> <p>The nodes in your chassis might look different depending on the type of nodes (storage or compute) you have.</p>
H610C chassis	
H615C chassis	
H610S chassis	

2. Unplug the power cord from the power supply unit.
3. Lift the cam handle, and press the blue latch to slide out the power supply unit.

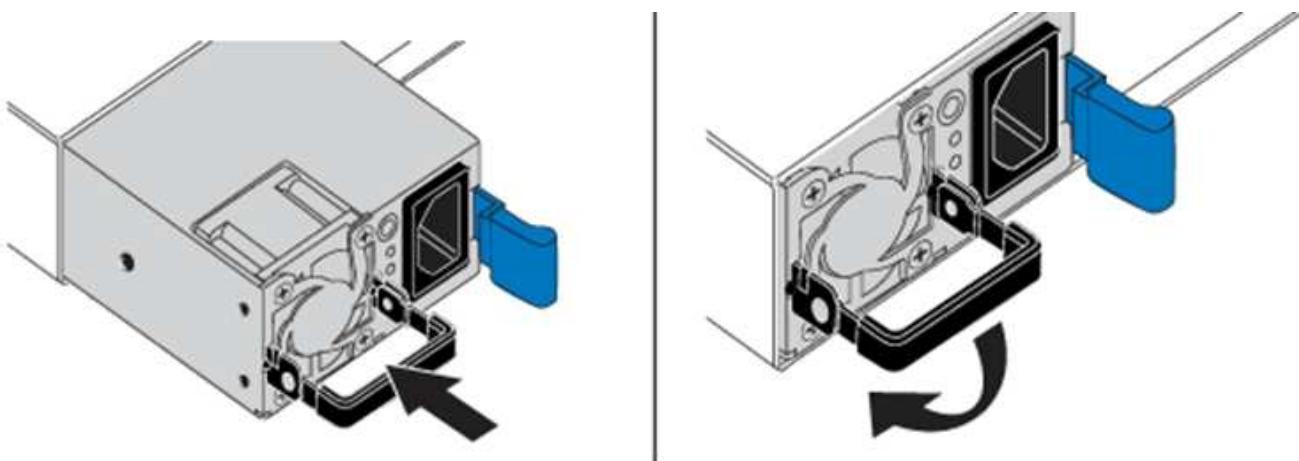


The illustration is an example. The location of the power supply unit in the chassis and the color of the release button vary depending on the type of chassis you have.



Ensure that you use both hands to support the weight of the power supply unit.

4. Using both hands, align the edges of the power supply unit with the opening in the chassis, gently push the unit into the chassis using the cam handle until it locks into place, and return the cam handle to the upright position.



5. Plug in the power cord.
6. Return the faulty unit to NetApp by following the instructions in the box that was shipped to you.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace SN2010, SN2100, and SN2700 switches

You can replace a faulty SN2000 series switch non-disruptively by following the best practices and steps provided by NetApp.

What you'll need

- Ensure that Putty is installed on the laptop and that you capture the output. See this video to learn how to configure Putty to capture the output session.

 | <https://img.youtube.com/vi/2LZfWH8HffA/maxresdefault.jpg>

- Ensure that you run NetApp Config Advisor before and after the replacement. This can help identify other problems before the maintenance starts. Download and install Config Advisor, and access the Quick Start Guide from [here \(login required\)](#).
- Obtain a power cable, the basic hand tools, and labels.
- Ensure that you have planned for a two to four-hour maintenance window.
- Familiarize yourself with the switch ports below:

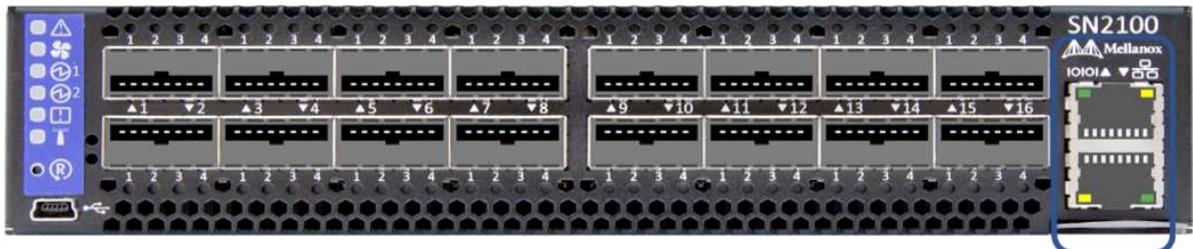


**18 Ports 1/10/25 GbE
SFP+ / SFP28**

**4 Ports 40/100 GbE
QSFP**

**Management
Ports**

Figure 1: SN2010 switch faceplate and ports



16 Ports 40/100 GbE, QSFP

Management Ports

Figure 2: SN2100 switch faceplate and ports

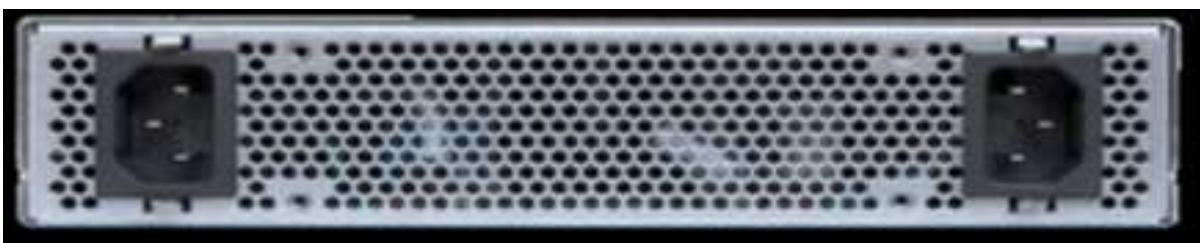


Figure 3: SN2010 and SN2100 switch rear

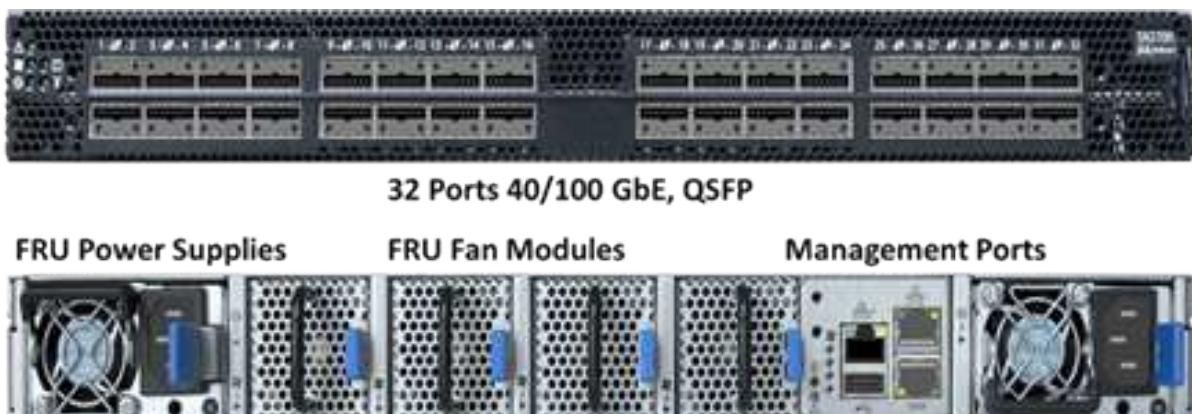


Figure 4: SN2700 switch front and rear

About this task

You should perform the steps in this procedure in the order below. This is to ensure that the downtime is minimal and the replacement switch is pre-configured before the switch replacement.



Contact NetApp Support if you need guidance.

Here is an overview of the steps in the procedure:

- [Prepare to replace the faulty switch](#)
- [Create the configuration file](#)
- [Remove the faulty switch and install the replacement](#)
- [Verify the operating system version on the switch](#)
- [Configure the replacement switch](#)
- [Complete the replacement](#)

Prepare to replace the faulty switch

Perform the following steps before you replace the faulty switch.

Steps

1. Verify that the replacement switch is the same model as the faulty switch.
2. Label all the cables connected to the faulty switch.
3. Identify the external file server where the switch configuration files are saved.
4. Ensure that you have obtained the following information:
 - a. The interface used for the initial configuration: RJ-45 port or the Serial Terminal Interface.
 - b. The credentials needed for switch access: IP address of the management port of the non-faulty switch and the faulty switch.
 - c. The passwords for administration access.

Create the configuration file

You can configure a switch by using the configuration files you create. Choose from one of the following options to create the configuration file for the switch.

Option	Steps
Create the backup configuration file from the faulty switch	<p>1. Connect to your switch remotely using SSH as shown in the following example:</p> <div data-bbox="897 295 1387 327" style="border: 1px solid black; padding: 5px; margin: 10px 0;"> <pre>ssh admin@<switch_IP_address></pre> </div> <p>2. Enter Configuration mode as shown in the following example:</p> <div data-bbox="897 528 1370 591" style="border: 1px solid black; padding: 5px; margin: 10px 0;"> <pre>switch > enable switch # configure terminal</pre> </div> <p>3. Find the available configuration files as shown in the following example:</p> <div data-bbox="897 802 1289 908" style="border: 1px solid black; padding: 5px; margin: 10px 0;"> <pre>switch (config) # switch (config) # show configuration files</pre> </div> <p>4. Save the active BIN configuration file to an external server:</p> <div data-bbox="897 1108 1444 1256" style="border: 1px solid black; padding: 5px; margin: 10px 0;"> <pre>switch (config) # configuration upload my-filename scp://myusername@my- server/path/to/my/<file></pre> </div>

Option	Steps
Create the backup configuration file by modifying the file from another switch	<p>1. Connect to your switch remotely using SSH as shown in the following example:</p> <pre data-bbox="899 297 1388 329">ssh admin@<switch_IP_address</pre> <p>2. Enter Configuration mode as shown in the following example:</p> <pre data-bbox="899 530 1372 593">switch > enable switch # configure terminal</pre> <p>3. Upload a text-based configuration file from the switch to an external server as shown in the following example:</p> <pre data-bbox="899 836 1445 1026">switch (config) # switch (config) # configuration text file my-filename upload scp://root@my- server/root/tmp/my-filename</pre> <p>4. Modify the following fields in the text file to match the faulty switch:</p> <pre data-bbox="899 1227 1413 1660"> ## Network interface configuration ## no interface mgmt0 dhcp interface mgmt0 ip address XX.XXX.XX.XXX /22 ## ## Other IP configuration ## hostname oldhostname</pre>

Remove the faulty switch and install the replacement

Perform the steps to remove the faulty switch and the install the replacement.

Steps

1. Locate the power cables on the faulty switch.
2. Label and unplug the power cables after the switch reboots.
3. Label and unplug all the cables from the faulty switch and secure them to prevent damage during switch replacement.
4. Remove the switch from the rack.
5. Install the replacement switch in the rack.
6. Connect the power cables and management port cables.



The switch automatically powers on when AC power is applied. There is no power button. It might take up to five minutes for the System Status LED to turn green.

7. Connect to the switch using the RJ-45 management port or the Serial Terminal Interface.

Verify the operating system version on the switch

Verify the OS software version on the switch. The version on the faulty switch and the healthy switch should match.

Steps

1. Connect to your switch remotely using SSH.
2. Enter Configuration mode.
3. Run the `show version` command. See the following example:

```
SFPS-HCI-SW02-A (config) #show version
Product name: Onyx
Product release: 3.7.1134
Build ID: #1-dev
Build date: 2019-01-24 13:38:57
Target arch: x86_64
Target hw: x86_64
Built by: jenkins@e4f385ab3f49
Version summary: x86_64 3.7.1134 2019-01-24 13:38:57 x86_64

Product model: x86onie
Host ID: 506B4B3238F8
System serial num: MT1812X24570
System UUID: 27fe4e7a-3277-11e8-8000-506b4b891c00

Uptime: 307d 3h 6m 33.344s
CPU load averages: 2.40 / 2.27 / 2.21
Number of CPUs: 4
System memory: 3525 MB used / 3840 MB free / 7365 MB total
Swap: 0 MB used / 0 MB free / 0 MB total
```

4. If the versions do not match, you should upgrade the OS. See the [Mellanox software Upgrade Guide](#) for

details.

Configure the replacement switch

Perform the steps to configure the replacement switch. See [Mellanox configuration management](#) for details.

Steps

1. Choose from the option that applies to you:

Option	Steps
From the BIN configuration file	<ol style="list-style-type: none">1. Fetch the BIN configuration file as shown in the following example:<pre>switch (config) # configuration fetch scp://myusername@my- server/path/to/my/<file></pre>2. Load the BIN configuration file you fetched in the previous step as shown in the following example:<pre>switch (config) # configuration switch-to my-filename</pre>3. Type <code>yes</code> to confirm the reboot.

Option	Steps
From the text file	<p>1. Reset the switch to factory default:</p> <pre data-bbox="897 253 1452 327">switch (config) # reset factory keep-basic</pre> <p>2. Apply the text-based configuration file:</p> <pre data-bbox="897 496 1452 570">switch (config) # configuration text file my-filename apply</pre> <p>3. Upload a text-based configuration file from the switch to an external server as shown in the following example:</p> <pre data-bbox="897 813 1452 1003">switch (config) # switch (config) # configuration text file my-filename upload scp://root@my- server/root/tmp/my-filename</pre> <p> A reboot is not required when you apply the text file.</p>

Complete the replacement

Perform the steps to complete the replacement procedure.

Steps

1. Insert the cables by using the labels to guide you.
2. Run NetApp Config Advisor. Access the Quick Start Guide from [here \(login required\)](#).
3. Verify your storage environment.
4. Return the faulty switch to NetApp.

Find more information

- [NetApp HCI Resources page](#)
- [SolidFire and Element Software Documentation Center](#)

Replace storage node in a two-node cluster

Before you replace a storage node that is part of a two-node cluster, you should first add a third storage node (which requires a new set of IP addresses), allow syncing to complete, and then remove your faulty node.

The cluster stays in the degraded state until a replacement node joins the cluster.

What you'll need

- You have new management IP and storage IP addresses.
- You have verified that the cluster shows the `ClusterCannotSync` alert after the node goes offline. This ensures that the cluster does a full resync when the new node is added back to the cluster. This alert is displayed approximately six minutes after the storage node goes offline.
- You have contacted NetApp Support.
If you are ordering a replacement, you should have a case open with NetApp Support.
- You have obtained the replacement node.
- You have an electrostatic discharge (ESD) wristband, or you have taken other antistatic protection.

About this task

Alarms in the VMware vSphere Web Client alert you when a host fails. You must match the serial number of the failed host from the VMware vSphere Web Client with the serial number on the sticker at the back of the node.

Steps

1. Physically remove the faulty node from the rack. The steps depend on the type of storage node you have. See [Replace H410S nodes](#) and [Replace H610S nodes](#).

 Do not remove the node from the cluster at this point.
2. Install the replacement node in the same slot.
3. Cable the node.
4. Power on the node.
5. Connect a keyboard and monitor to the node.
6. Perform the configuration steps:
 - a. Configure the IPMI/BMC IP address.
 - b. Configure the new node with the new management IP and storage IP addresses, and the Cluster Name.
7. After the node is added to the cluster, add the drives.
8. After the sync is finished, remove the failed drives and the failed node from the cluster.
9. Use NetApp Hybrid Cloud Control to configure the new storage node that you added. See [Expand NetApp HCI storage resources](#).

Find more information

- [NetApp HCI Documentation Center](#)
- [SolidFire and Element Software Documentation Center](#)

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

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.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

- [Notice for Compute Firmware Bundle 2.146](#)
- [Notice for Storage Firmware Bundle 2.146](#)
- [Notice for Storage Firmware Bundle 2.99.2](#)
- [Notice for Compute Firmware Bundle 2.76](#)
- [Notice for Storage Firmware Bundle 2.76](#)
- [Notice for Compute Firmware Bundle 2.27](#)
- [Notice for Storage Firmware Bundle 2.27](#)
- [Notice for compute firmware ISO](#)
- [Notice for H610S BMC](#)
- [Notice for Management Services 2.23.64 \(NetApp Element Plug-in for VMware vCenter Server 5.1.12\)](#)
- [Notice for Management Services 2.22.7 \(NetApp Element Plug-in for VMware vCenter Server 5.0.37\)](#)
- [Notice for Management Services 2.21.61 \(NetApp Element Plug-in for vCenter Server 4.10.12\)](#)
- [Notice for Management Services 2.20.69 \(NetApp Element Plug-in for vCenter Server 4.9.14\)](#)
- [Notice for Management Services 2.19.48 \(NetApp Element Plug-in for vCenter Server 4.8.34\)](#)
- [Notice for Management Services 2.18.91 \(NetApp Element Plug-in for vCenter Server 4.7.10\)](#)

- Notice for Management Services 2.17.56 (NetApp Element Plug-in for vCenter Server 4.6.32)
- Notice for Management Services 2.17 (NetApp Element Plug-in for vCenter Server 4.6.29)
- Notice for Management Services 2.16 (NetApp Element Plug-in for vCenter Server 4.6.29)
- Notice for Management Services 2.14 (NetApp Element Plug-in for vCenter Server 4.5.42)
- Notice for Management Services 2.13 (NetApp Element Plug-in for vCenter Server 4.5.42)
- Notice for Management Services 2.11 (NetApp Element Plug-in for vCenter Server 4.4.72)
- Notice for NetApp HCI 1.8

Copyright information

Copyright © 2025 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—with 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.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, 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) (FEB 2014).

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.