



Troubleshoot your system

Element Software

NetApp

November 18, 2025

This PDF was generated from https://docs.netapp.com/us-en/element-software-128/storage/task_monitor_information_about_system_events.html on November 18, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Troubleshoot your system	1
System events	1
View information about system events	1
Event types	1
View status of running tasks	5
System alerts	5
View system alerts	5
Cluster fault codes	6
View node performance activity	22
Volume performance	22
View volume performance	22
Volume performance details	23
iSCSI sessions	24
View iSCSI sessions	24
iSCSI session details	25
Fibre Channel sessions	26
View Fibre Channel sessions	26
Fibre Channel session details	26
Troubleshoot drives	27
Troubleshoot drives	27
Remove failed drives from the cluster	28
Basic MDSS drive troubleshooting	28
Add MDSS drives	29
Remove MDSS drives	30
Troubleshoot nodes	30
Remove nodes from a cluster	30
Power down a cluster	31
Work with per-node utilities for storage nodes	31
Work with per-node utilities for storage nodes	31
Access per-node settings using the per-node UI	32
Network settings details from the per-node UI	33
Cluster settings details from the per-node UI	35
Run system tests using the per-node UI	36
Run system utilities using the per-node UI	37
Work with the management node	39
Understand cluster fullness levels	39

Troubleshoot your system

System events

View information about system events

You can view information about various events detected in the system. The system refreshes the event messages every 30 seconds. The event log displays key events for the cluster.

1. In the Element UI, select **Reporting > Event Log**.

For every event, you see the following information:

Item	Description
ID	Unique ID associated with each event.
Event Type	The type of event being logged, for example, API events or clone events.
Message	Message associated with the event.
Details	Information that helps identify why the event occurred.
Service ID	The service that reported the event (if applicable).
Node	The node that reported the event (if applicable).
Drive ID	The drive that reported the event (if applicable).
Event Time	The time the event occurred.

Find more information

[Event types](#)

Event types

The system reports multiple types of events; each event is an operation that the system has completed. Events can be routine, normal events or events that require administrator attention. The Event Types column on the Event Log page indicates in which part of the system the event occurred.



The system does not log read-only API commands in the event log.

The following list describes the types of events that appear in the event log:

- **apiEvent**

Events initiated by a user through an API or web UI that modify settings.

- **binAssignmentsEvent**

Events related to the assignment of data bins. Bins are essentially containers that hold data and are mapped across the cluster.

- **binSyncEvent**

System events related to a reassignment of data among block services.

- **bsCheckEvent**

System events related to block service checks.

- **bsKillEvent**

System events related to block service terminations.

- **bulkOpEvent**

Events related to operations performed on an entire volume, such as a backup, restore, snapshot, or clone.

- **cloneEvent**

Events related to volume cloning.

- **clusterMasterEvent**

Events appearing upon cluster initialization or upon configuration changes to the cluster, such as adding or removing nodes.

- **cSumEvent**

Events related to the detection of a checksum mismatch during end-to-end checksum validation.

Services that detect a checksum mismatch are automatically stopped and not restarted after generating this event.

- **dataEvent**

Events related to reading and writing data.

- **dbEvent**

Events related to the global database maintained by ensemble nodes in the cluster.

- **driveEvent**

Events related to drive operations.

- **encryptionAtRestEvent**

Events related to the process of encryption on a cluster.

- **ensembleEvent**

Events related to increasing or decreasing the number of nodes in an ensemble.

- **fibreChannelEvent**

Events related to the configuration of and connections to the nodes.

- **gcEvent**

Events related to processes run every 60 minutes to reclaim storage on block drives. This process is also known as garbage collection.

- **ieEvent**

Internal system error.

- **installEvent**

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

- **iSCSIEvent**

Events related to iSCSI issues in the system.

- **limitEvent**

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

- **maintenanceModeEvent**

Events related to the node maintenance mode, such as disabling the node.

- **networkEvent**

Events related to the network error reporting for each physical network interface card (NIC) interface.

These events are triggered when any error count for an interface exceeds a default threshold of 1000 during a 10-minute monitoring interval. These events apply to network errors such as received misses, cyclic redundancy check (CRC) errors, length errors, overrun errors, and frame errors.

- **platformHardwareEvent**

Events related to issues detected on hardware devices.

- **remoteClusterEvent**

Events related to remote cluster pairing.

- **schedulerEvent**

Events related to scheduled snapshots.

- **serviceEvent**

Events related to system service status.

- **sliceEvent**

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

There are three types of slice reassignment events, which include information about the service where a volume is assigned:

- flipping: changing the primary service to a new primary service

```
sliceID oldPrimaryServiceID->newPrimaryServiceID
```

- moving: changing the secondary service to a new secondary service

```
sliceID {oldSecondaryServiceID(s)}->{newSecondaryServiceID(s)}
```

- pruning: removing a volume from a set of services

```
sliceID {oldSecondaryServiceID(s)}
```

- **snmpTrapEvent**

Events related to SNMP traps.

- **statEvent**

Events related to system statistics.

- **tsEvent**

Events related to the system transport service.

- **unexpectedException**

Events related to unexpected system exceptions.

- **ureEvent**

Events related to Unrecoverable Read Errors that occur while reading from the storage device.

- **vasaProviderEvent**

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

View status of running tasks

You can view the progress and completion status of running tasks in the web UI that are being reported by the ListSyncJobs and ListBulkVolumeJobs API methods. You can access the Running Tasks page from the Reporting tab of the Element UI.

If there are a large number of tasks, the system might queue them and run them in batches. The Running Tasks page displays the services currently being synchronized. When a task is complete, it is replaced by the next queued synchronizing task. Synchronizing tasks might continue to appear on the Running Tasks page until there are no more tasks to complete.



You can see replication synchronizations data for volumes undergoing replication on the Running Tasks page of the cluster containing the target volume.

System alerts

View system alerts

You can view alerts for information about cluster faults or errors in the system. Alerts can be information, warnings, or errors and are a good indicator of how well the cluster is running. Most errors resolve themselves automatically.

You can use the ListClusterFaults API method to automate alert monitoring. This enables you to be notified about all alerts that occur.

1. In the Element UI, select **Reporting > Alerts**.

The system refreshes the alerts on the page every 30 seconds.

For every event, you see the following information:

Item	Description
ID	Unique ID associated with a cluster alert.

Severity	The degree of importance of the alert. Possible values: • warning: A minor issue that might soon require attention. System upgrades are still allowed. • error: A failure that might cause performance degradation or loss of high availability (HA). Errors generally should not affect service otherwise. • critical: A serious failure that affects service. The system is unable to serve API or client I/O requests. Operating in this state could lead to potential loss of data. • bestPractice: A recommended system configuration best practice is not being used.
Type	The component that the fault affects. Can be node, drive, cluster, service, or volume.
Node	Node ID for the node that this fault refers to. Included for node and drive faults, otherwise set to - (dash).
Drive ID	Drive ID for the drive that this fault refers to. Included for drive faults, otherwise set to - (dash).
Error Code	A descriptive code that indicates what caused the fault.
Details	A description of the fault with additional details.
Date	The date and time the fault was logged.

2. Click **Show Details** for an individual alert to view information about the alert.

3. To view the details of all alerts on the page, click the Details column.

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

Find more information

- [Cluster fault codes](#)
- [Manage storage with the Element API](#)

Cluster fault codes

The system reports an error or a state that might be of interest by generating a fault code,

which is listed on the Alerts page. These codes help you determine what component of the system experienced the alert and why the alert was generated.

The following list outlines the different types of codes:

- **authenticationServiceFault**

The Authentication Service on one or more cluster nodes is not functioning as expected.

Contact NetApp Support for assistance.

- **availableVirtualNetworkIPAddressesLow**

The number of virtual network addresses in the block of IP addresses is low.

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

- **blockClusterFull**

There is not enough free block storage space to support a single node loss. See the `GetClusterFullThreshold` API method for details on cluster fullness levels. This cluster fault indicates one of the following conditions:

- `stage3Low` (Warning): User-defined threshold was crossed. Adjust Cluster Full settings or add more nodes.
- `stage4Critical` (Error): There is not enough space to recover from a 1-node failure. Creation of volumes, snapshots, and clones is not allowed.
- `stage5CompletelyConsumed` (Critical)¹: No writes or new iSCSI connections are allowed. Current iSCSI connections will be maintained. Writes will fail until more capacity is added to the cluster.

To resolve this fault, purge or delete volumes or add another storage node to the storage cluster.

- **blocksDegraded**

Block data is no longer fully replicated due to a failure.

Severity	Description
Warning	Only two complete copies of the block data are accessible.
Error	Only a single complete copy of the block data is accessible.
Critical	No complete copies of the block data are accessible.

Note: The warning status can only occur on a Triple Helix system.

To resolve this fault, restore any offline nodes or block services, or contact NetApp Support for assistance.

- **blockServiceTooFull**

A block service is using too much space.

To resolve this fault, add more provisioned capacity.

- **blockServiceUnhealthy**

A block service has been detected as unhealthy:

- Severity = Warning: No action is taken. This warning period will expire in `cTimeUntilBSIsKilledMSec=330000` milliseconds.
- Severity = Error: The system is automatically decommissioning data and re-replicating its data to other healthy drives.
- Severity = Critical: There are failed block services on several nodes greater than or equal to the replication count (2 for double helix). Data is unavailable and bin syncing will not finish.

Check for network connectivity issues and hardware errors. There will be other faults if specific hardware components have failed. The fault will clear when the block service is accessible or when the service has been decommissioned.

- **BmcSelfTestFailed**

The Baseboard Management Controller (BMC) failed a self-test.

Contact NetApp support for assistance.

During an upgrade to Element 12.5 or later, the `BmcSelfTestFailed` fault is not generated for a node that has a preexisting failed BMC, or when a node's BMC fails during the upgrade. The BMCs that fail the self-tests during the upgrade will issue a `BmcSelfTestFailed` warning fault after the entire cluster completes the upgrade.

- **clockSkewExceedsFaultThreshold**

Time skew between the Cluster master and the node which is presenting a token exceeds the recommended threshold. Storage cluster cannot correct the time skew between the nodes automatically.

To resolve this fault, use NTP servers that are internal to your network, rather than the installation defaults. If you are using an internal NTP server, contact NetApp Support for assistance.

- **clusterCannotSync**

There is an out-of-space condition and data on the offline block storage drives cannot be synced to drives that are still active.

To resolve this fault, add more storage.

- **clusterFull**

There is no more free storage space in the storage cluster.

To resolve this fault, add more storage.

- **clusterIOPSAreOverProvisioned**

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

To resolve this issue, lower the minimum QoS IOPS settings for volumes.

- **CpuThermalEventThreshold**

The number of CPU thermal events on one or more CPUs exceeds the configured threshold.

If no new CPU thermal events are detected within ten minutes, the warning will resolve itself.

- **disableDriveSecurityFailed**

The cluster is not configured to enable drive security (Encryption at Rest), but at least one drive has drive security enabled, meaning that disabling drive security on those drives failed. This fault is logged with “Warning” severity.

To resolve this fault, check the fault details for the reason why drive security could not be disabled. Possible reasons are:

- The encryption key could not be acquired, investigate the problem with access to the key or the external key server.
- The disable operation failed on the drive, determine whether the wrong key could possibly have been acquired.

If neither of these are the reason for the fault, the drive might need to be replaced.

You can attempt to recover a drive that does not successfully disable security even when the correct authentication key is provided. To perform this operation, remove the drive(s) from the system by moving it to Available, perform a secure erase on the drive and move it back to Active.

- **disconnectedClusterPair**

A cluster pair is disconnected or configured incorrectly.

Check network connectivity between the clusters.

- **disconnectedRemoteNode**

A remote node is either disconnected or configured incorrectly.

Check network connectivity between the nodes.

- **disconnectedSnapMirrorEndpoint**

A remote SnapMirror endpoint is disconnected or configured incorrectly.

Check network connectivity between the cluster and the remote SnapMirrorEndpoint.

- **driveAvailable**

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

To resolve this fault, add any available drives to the storage cluster.

- **driveFailed**

The cluster returns this fault when one or more drives have failed, indicating one of the following

conditions:

- The drive manager cannot access the drive.
- The slice or block service has failed too many times, presumably because of drive read or write failures, and cannot restart.
- The drive is missing.
- The master service for the node is inaccessible (all drives in the node are considered missing/failed).
- The drive is locked and the authentication key for the drive cannot be acquired.
- The drive is locked and the unlock operation fails.

To resolve this issue:

- Check network connectivity for the node.
- Replace the drive.
- Ensure that the authentication key is available.

• **driveHealthFault**

A drive has failed the SMART health check and as a result, the drive's functions are diminished. There is a Critical severity level for this fault:

- 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.

• **driveWearFault**

A drive's remaining life has dropped below thresholds, but it is still functioning. There are two possible severity levels for this fault: Critical and Warning:

- Drive with serial: <serial number> in slot: <node slot><drive slot> has critical wear levels.
- Drive with serial: <serial number> in slot: <node slot><drive slot> has low wear reserves.

To resolve this fault, replace the drive soon.

• **duplicateClusterMasterCandidates**

More than one storage cluster master candidate has been detected.

Contact NetApp Support for assistance.

• **enableDriveSecurityFailed**

The cluster is configured to require drive security (Encryption at Rest), but drive security could not be enabled on at least one drive. This fault is logged with "Warning" severity.

To resolve this fault, check the fault details for the reason why drive security could not be enabled. Possible reasons are:

- The encryption key could not be acquired, investigate the problem with access to the key or the external key server.

- The enable operation failed on the drive, determine whether the wrong key could possibly have been acquired. If neither of these are the reason for the fault, the drive might need to be replaced.

You can attempt to recover a drive that does not successfully enable security even when the correct authentication key is provided. To perform this operation, remove the drive(s) from the system by moving it to Available, perform a secure erase on the drive and move it back to Active.

- **ensembleDegraded**

Network connectivity or power has been lost to one or more of the ensemble nodes.

To resolve this fault, restore network connectivity or power.

- **exception**

A fault reported that is other than a routine fault. These faults are not automatically cleared from the fault queue.

Contact NetApp Support for assistance.

- **failedSpaceTooFull**

A block service is not responding to data write requests. This causes the slice service to run out of space to store failed writes.

To resolve this fault, restore block services functionality to allow writes to continue normally and failed space to be flushed from the slice service.

- **fanSensor**

A fan sensor has failed or is missing.

To resolve this fault, replace any failed hardware.

- **fibreChannelAccessDegraded**

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

Check network connectivity.

- **fibreChannelAccessUnavailable**

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

Check network connectivity.

- **fibreChannelActiveIxl**

The IxL Nexus count is approaching the supported limit of 8000 active sessions per Fibre Channel node.

- Best practice limit is 5500.
- Warning limit is 7500.
- Maximum limit (not enforced) is 8192.

To resolve this fault, reduce the IxL Nexus count below the best practice limit of 5500.

- **fibreChannelConfig**

This cluster fault indicates one of the following conditions:

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

Contact NetApp Support for assistance.

- **fibreChannelIOPS**

The total IOPS count is approaching the IOPS limit for Fibre Channel nodes in the cluster. The limits are:

- FC0025: 450K IOPS limit at 4K block size per Fibre Channel node.
- FCN001: 625K OPS limit at 4K block size per Fibre Channel node.

To resolve this fault, balance the load across all available Fibre Channel nodes.

- **fibreChannelStaticIxl**

The Ixl Nexus count is approaching the supported limit of 16000 static sessions per Fibre Channel node.

- Best practice limit is 11000.
- Warning limit is 15000.
- Maximum limit (enforced) is 16384.

To resolve this fault, reduce the Ixl Nexus count below the best practice limit of 11000.

- **fileSystemCapacityLow**

There is insufficient space on one of the filesystems.

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

- **fileSystemIsReadOnly**

A filesystem has moved into read-only mode.

Contact NetApp Support for assistance.

- **fipsDrivesMismatch**

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

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

- **fipsDrivesOutOfCompliance**

The system has detected that Encryption at Rest was disabled after the FIPS Drives feature was enabled.

This fault is also generated when the FIPS Drives feature is enabled and a non-FIPS drive or node is present in the storage cluster.

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

- **fipsSelfTestFailure**

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

Contact NetApp Support for assistance.

- **hardwareConfigMismatch**

This cluster fault indicates one of the following conditions:

- The configuration does not match the node definition.
- There is an incorrect drive size for this type of node.
- An unsupported drive has been detected. A possible reason is that the installed Element version does not recognize this drive. Recommend updating the Element software on this node.
- There is a drive firmware mismatch.
- The drive encryption capable state does not match the node.

Contact NetApp Support for assistance.

- **idPCertificateExpiration**

The cluster's service provider SSL certificate for use with a third-party identity provider (IdP) is nearing expiration or has already expired. This fault uses the following severities based on urgency:

Severity	Description
Warning	Certificate expires within 30 days.
Error	Certificate expires within 7 days.
Critical	Certificate expires within 3 days or has already expired.

To resolve this fault, update the SSL certificate before it expires. Use the `UpdateIdpConfiguration` API method with `refreshCertificateExpirationTime=true` to provide the updated SSL certificate.

- **inconsistentBondModes**

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

- **inconsistentMtus**

This cluster fault indicates one of the following conditions:

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

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

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

- **inconsistentRoutingRules**

The routing rules for this interface are inconsistent.

- **inconsistentSubnetMasks**

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

- **incorrectBondPortCount**

The number of bond ports is incorrect.

- **invalidConfiguredFibreChannelNodeCount**

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

To resolve this fault, check the cluster network connectivity and network cabling, and check for failed services. If there are no network or service problems, contact NetApp Support for a Fibre Channel node replacement.

- **irqBalanceFailed**

An exception occurred while attempting to balance interrupts.

Contact NetApp Support for assistance.

- **kmipCertificateFault**

- Root Certification Authority (CA) certificate is nearing expiration.

To resolve this fault, acquire a new certificate from the root CA with expiration date at least 30 days out and use `ModifyKeyServerKmip` to provide the updated root CA certificate.

- Client certificate is nearing expiration.

To resolve this fault, create a new CSR using `GetClientCertificateSigningRequest`, have it signed ensuring the new expiration date is at least 30 days out, and use `ModifyKeyServerKmip` to replace the expiring KMIP client certificate with the new certificate.

- Root Certification Authority (CA) certificate has expired.

To resolve this fault, acquire a new certificate from the root CA with expiration date at least 30 days out and use `ModifyKeyServerKmip` to provide the updated root CA certificate.

- Client certificate has expired.

To resolve this fault, create a new CSR using `GetClientCertificateSigningRequest`, have it signed ensuring the new expiration date is at least 30 days out, and use `ModifyKeyServerKmip` to replace the expired KMIP client certificate with the new certificate.

- Root Certification Authority (CA) certificate error.

To resolve this fault, check that the correct certificate was provided, and, if needed, reacquire the certificate from the root CA. Use `ModifyKeyServerKmip` to install the correct KMIP client certificate.

- Client certificate error.

To resolve this fault, check that the correct KMIP client certificate is installed. The root CA of the client certificate should be installed on the EKS. Use `ModifyKeyServerKmip` to install the correct KMIP client certificate.

- **kmipServerFault**

- Connection failure

To resolve this fault, check that the External Key Server is alive and reachable via the network. Use `TestKeyServerKimp` and `TestKeyProviderKmip` to test your connection.

- Authentication failure

To resolve this fault, check that the correct root CA and KMIP client certificates are being used, and that the private key and the KMIP client certificate match.

- Server error

To resolve this fault, check the details for the error. Troubleshooting on the External Key Server might be necessary based on the error returned.

- **memoryEccThreshold**

A large number of correctable or uncorrectable ECC errors have been detected. This fault uses the following severities based on urgency:

Event	Severity	Description
A single DIMM <code>cErrorCount</code> reaches <code>cDimmCorrectableErrWarnThreshold</code> .	Warning	Correctable ECC memory errors above threshold on DIMM: <Processor> <DIMM Slot>
A single DIMM <code>cErrorCount</code> stays above <code>cDimmCorrectableErrWarnThreshold</code> until <code>cErrorFaultTimer</code> expires for the DIMM.	Error	Correctable ECC memory errors above threshold on DIMM: <Processor> <DIMM>
A memory controller reports <code>cErrorCount</code> above <code>cMemCtrlCorrectableErrWarnThreshold</code> , and <code>cMemCtrlCorrectableErrWarnDuration</code> is specified.	Warning	Correctable ECC memory errors above threshold on memory controller: <Processor> <Memory Controller>

A memory controller reports cErrorCount above cMemCtrlCorrectableErrWarnThreshold until cErrorFaultTimer expires for the memory controller.	Error	Correctable ECC memory errors above threshold on DIMM: <Processor> <DIMM>
A single DIMM reports a uErrorCount above zero, but less than cDimmUncorrectableErrFaultThreshold.	Warning	Uncorrectable ECC memory error(s) detected on DIMM: <Processor> <DIMM Slot>
A single DIMM reports a uErrorCount of at least cDimmUncorrectableErrFaultThreshold.	Error	Uncorrectable ECC memory error(s) detected on DIMM: <Processor> <DIMM Slot>
A memory controller reports a uErrorCount above zero, but less than cMemCtrlUncorrectableErrFaultThreshold.	Warning	Uncorrectable ECC memory error(s) detected on memory controller: <Processor> <Memory Controller>
A memory controller reports a uErrorCount of at least cMemCtrlUncorrectableErrFaultThreshold.	Error	Uncorrectable ECC memory error(s) detected on memory controller: <Processor> <Memory Controller>

To resolve this fault, contact NetApp Support for assistance.

- **memoryUsageThreshold**

Memory usage is above normal. This fault uses the following severities based on urgency:



See the **Details** heading in the error fault for more detailed information on the type of fault.

Severity	Description
Warning	System memory is low.
Error	System memory is very low.
Critical	System memory is completely consumed.

To resolve this fault, contact NetApp Support for assistance.

- **metadataClusterFull**

There is not enough free metadata storage space to support a single node loss. See the GetClusterFullThreshold API method for details on cluster fullness levels. This cluster fault indicates one of

the following conditions:

- **stage3Low (Warning)**: User-defined threshold was crossed. Adjust Cluster Full settings or add more nodes.
- **stage4Critical (Error)**: There is not enough space to recover from a 1-node failure. Creation of volumes, snapshots, and clones is not allowed.
- **stage5CompletelyConsumed (Critical)**1; No writes or new iSCSI connections are allowed. Current iSCSI connections will be maintained. Writes will fail until more capacity is added to the cluster. Purge or delete data or add more nodes.

To resolve this fault, purge or delete volumes or add another storage node to the storage cluster.

- **mtuCheckFailure**

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

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

- **networkConfig**

This cluster fault indicates one of the following conditions:

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

Contact NetApp Support for assistance.

- **noAvailableVirtualNetworkIPAddresses**

There are no available virtual network addresses in the block of IP addresses.

- **virtualNetworkID # TAG(###)** has no available storage IP addresses. Additional nodes cannot be added to the cluster.

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

- **nodeHardwareFault (Network interface <name> is down or cable is unplugged)**

A network interface is either down or the cable is unplugged.

To resolve this fault, check network connectivity for the node or nodes.

- **nodeHardwareFault (Drive encryption capable state mismatches node's encryption capable state for the drive in slot <node slot><drive slot>)**

A drive does not match encryption capabilities with the storage node it is installed in.

- **nodeHardwareFault (Incorrect <drive type> drive size <actual size> for the drive in slot <node slot><drive slot> for this node type - expected <expected size>)**

A storage node contains a drive that is the incorrect size for this node.

- **nodeHardwareFault (Unsupported drive detected in slot <node slot><drive slot>; drive statistics and health information will be unavailable)**

A storage node contains a drive it does not support.

- **nodeHardwareFault (The drive in slot <node slot><drive slot> should be using firmware version <expected version>, but is using unsupported version <actual version>)**

A storage node contains a drive running an unsupported firmware version.

- **nodeMaintenanceMode**

A node has been placed in maintenance mode. This fault uses the following severities based on urgency:

Severity	Description
Warning	Indicates that the node is still in maintenance mode.
Error	Indicates that maintenance mode has failed to disable, most likely due to failed or active standbys.

To resolve this fault, disable maintenance mode once maintenance completes. If the Error level fault persists, contact NetApp Support for assistance.

- **nodeOffline**

Element software cannot communicate with the specified node. Check network connectivity.

- **notUsingLACPBondMode**

LACP bonding mode is not configured.

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

- **ntpServerUnreachable**

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

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

- **ntpTimeNotInSync**

The difference between storage cluster time and the specified NTP server time is too large. The storage cluster cannot correct the difference automatically.

To resolve this fault, use NTP servers that are internal to your network, rather than the installation defaults. If you are using internal NTP servers and the issue persists, contact NetApp Support for assistance.

- **nvramDeviceStatus**

An NVRAM device has an error, is failing, or has failed. This fault has the following severities:

Severity	Description
----------	-------------

Warning	A warning has been detected by the hardware. This condition may be transitory, such as a temperature warning. • nvmLifetimeError • nvmLifetimeStatus • energySourceLifetimeStatus • energySourceTemperatureStatus • warningThresholdExceeded
Error	An Error or Critical status has been detected by the hardware. The cluster master attempts to remove the slice drive from operation (this generates a drive removal event). If secondary slice services are not available the drive will not be removed. Errors returned in addition to the Warning level errors: • NVRAM device mount point doesn't exist. • NVRAM device partition doesn't exist. • NVRAM device partition exists, but not mounted.
Critical	An Error or Critical status has been detected by the hardware. The cluster master attempts to remove the slice drive from operation (this generates a drive removal event). If secondary slice services are not available the drive will not be removed. • persistenceLost • armStatusSaveNArmed • csaveStatusError

Replace any failed hardware in the node. If this does not resolve the issue, contact NetApp Support for assistance.

- **powerSupplyError**

This cluster fault indicates one of the following conditions:

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

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

- **provisionedSpaceTooFull**

The overall provisioned capacity of the cluster is too full.

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

- **remoteRepAsyncDelayExceeded**

The configured asynchronous delay for replication has been exceeded. Check network connectivity between clusters.

- **remoteRepClusterFull**

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

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

- **remoteRepSnapshotClusterFull**

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

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

- **remoteRepSnapshotsExceededLimit**

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

To resolve this fault, increase the snapshot limit on the target storage cluster.

- **scheduleActionError**

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

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

- **sensorReadingFailed**

A sensor could not communicate with the Baseboard Management Controller (BMC).

Contact NetApp Support for assistance.

- **serviceNotRunning**

A required service is not running.

Contact NetApp Support for assistance.

- **sliceServiceTooFull**

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

To resolve this fault, add more provisioned capacity.

- **sliceServiceUnhealthy**

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

- Severity = Warning: No action is taken. This warning period will expire in 6 minutes.

- Severity = Error: The system is automatically decommissioning data and re-replicating its data to other healthy drives.

Check for network connectivity issues and hardware errors. There will be other faults if specific hardware components have failed. The fault will clear when the slice service is accessible or when the service has been decommissioned.

- **sshEnabled**

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

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

- **sslCertificateExpiration**

The SSL certificate associated with this node is nearing expiration or has expired. This fault uses the following severities based on urgency:

Severity	Description
Warning	Certificate expires within 30 days.
Error	Certificate expires within 7 days.
Critical	Certificate expires within 3 days or has already expired.

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

- **strandedCapacity**

A single node accounts for more than half of the storage cluster capacity.

In order to maintain data redundancy, the system reduces the capacity of the largest node so that some of its block capacity is stranded (not used).

To resolve this fault, add more drives to existing storage nodes or add storage nodes to the cluster.

- **tempSensor**

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

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

- **upgrade**

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

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

- **unresponsiveService**

A service has become unresponsive.

Contact NetApp Support for assistance.

- **virtualNetworkConfig**

This cluster fault indicates one of the following conditions:

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

Contact NetApp Support for assistance.

- **volumesDegraded**

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

- **volumesOffline**

One or more volumes in the storage cluster are offline. The **volumeDegraded** fault will also be present.

Contact NetApp Support for assistance.

View node performance activity

You can view performance activity for each node in a graphical format. This information provides real-time statistics for CPU and read/write I/O operations per second (IOPS) for each drive the node. The utilization graph is updated every five seconds, and the drive statistics graph updates every ten seconds.

1. Click **Cluster > Nodes**.
2. Click **Actions** for the node you want to view.
3. Click **View Details**.



You can see specific points in time on the line and bar graphs by positioning your cursor over the line or bar.

Volume performance

View volume performance

You can view detailed performance information for all volumes in the cluster. You can sort the information by volume ID or by any of the performance columns. You can also use

filter the information by certain criteria.

You can change how often the system refreshes performance information on the page by clicking the **Refresh every** list, and choosing a different value. The default refresh interval is 10 seconds if the cluster has less than 1000 volumes; otherwise, the default is 60 seconds. If you choose a value of Never, automatic page refreshing is disabled.

You can reenable automatic refreshing by clicking **Turn on auto-refresh**.

1. In the Element UI, select **Reporting > Volume Performance**.
2. In the volume list, click the Actions icon for a volume.
3. Click **View Details**.

A tray is displayed at the bottom of the page containing general information about the volume.

4. To see more detailed information about the volume, click **See More Details**.

The system displays detailed information as well as performance graphs for the volume.

Find more information

[Volume performance details](#)

Volume performance details

You can view performance statistics of volumes from the Volume Performance page of the Reporting tab in the Element UI.

The following list describes the details that are available to you:

- **ID**

The system-generated ID for the volume.

- **Name**

The name given to the volume when it was created.

- **Account**

The name of the account assigned to the volume.

- **Access Groups**

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

- **Volume Utilization**

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

Possible values:

- 0 = Client is not using the volume

- 100 = Client is using the max
- >100 = Client is using the burst

- **Total IOPS**

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

- **Read IOPS**

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

- **Write IOPS**

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

- **Total Throughput**

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

- **Read Throughput**

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

- **Write Throughput**

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

- **Total Latency**

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

- **Read Latency**

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

- **Write Latency**

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

- **Queue Depth**

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

- **Average IO Size**

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

iSCSI sessions

View iSCSI sessions

You can view the iSCSI sessions that are connected to the cluster. You can filter the information to include only the desired sessions.

1. In the Element UI, select **Reporting > iSCSI Sessions**.
2. To see the filter criteria fields, click **Filter**.

Find more information

[iSCSI session details](#)

iSCSI session details

You can view information about the iSCSI sessions that are connected to the cluster.

The following list describes the information that you can find about the iSCSI sessions:

- **Node**

The node hosting the primary metadata partition for the volume.

- **Account**

The name of the account that owns the volume. If value is blank, a dash (-) is displayed.

- **Volume**

The volume name identified on the node.

- **Volume ID**

ID of the volume associated with the Target IQN.

- **Initiator ID**

A system-generated ID for the initiator.

- **Initiator Alias**

An optional name for the initiator that makes finding the initiator easier when in a long list.

- **Initiator IP**

The IP address of the endpoint that initiates the session.

- **Initiator IQN**

The IQN of the endpoint that initiates the session.

- **Target IP**

The IP address of the node hosting the volume.

- **Target IQN**

The IQN of the volume.

- **CHAP**

The CHAP algorithm for an iSCSI session. If a CHAP algorithm is not being used, a dash (-) is displayed. Available beginning with Element 12.8.

- **Created On**

Date the session was established.

Fibre Channel sessions

View Fibre Channel sessions

You can view the Fibre Channel (FC) sessions that are connected to the cluster. You can filter information to include only those connections you want displayed in the window.

1. In the Element UI, select **Reporting > FC Sessions**.
2. To see the filter criteria fields, click **Filter**.

Find more information

[Fibre Channel session details](#)

Fibre Channel session details

You can find information about the active Fibre Channel (FC) sessions that are connected to the cluster.

The following list describes the information you can find about the FC sessions connected to the cluster:

- **Node ID**

The node hosting the session for the connection.

- **Node Name**

System-generated node name.

- **Initiator ID**

A system-generated ID for the initiator.

- **Initiator WWPN**

The initiating worldwide port name.

- **Initiator Alias**

An optional name for the initiator that makes finding the initiator easier when in a long list.

- **Target WWPN**

The target worldwide port name.

- **Volume Access Group**

Name of the volume access group that the session belongs to.

- **Volume Access Group ID**

System-generated ID for the access group.

Troubleshoot drives

Troubleshoot drives

You can replace a failed solid-state drive (SSD) with a replacement drive. SSDs for SolidFire storage nodes are hot-swappable. If you suspect an SSD has failed, contact NetApp Support to verify the failure and walk you through the proper resolution procedure. NetApp Support also works with you to get a replacement drive according to your service-level agreement.

How-swappable in this case means that you can remove a failed drive from an active node and replace it with a new SSD drive from NetApp. It is not recommended that you should remove non-failed drives on an active cluster.

You should maintain on-site spares suggested by NetApp Support to allow for immediate replacement of the drive if it fails.



For testing purposes, if you are simulating a drive failure by pulling a drive from a node, you must wait 30 seconds before inserting the drive back into the drive slot.

If a drive fails, Double Helix redistributes the data on the drive across the nodes remaining on the cluster. Multiple drive failures on the same node are not an issue since Element software protects against two copies of data residing on the same node. A failed drive results in the following events:

- Data is migrated off of the drive.
- Overall cluster capacity is reduced by the capacity of the drive.
- Double Helix data protection ensures that there are two valid copies of the data.



SolidFire storage systems do not support removal of a drive if it results in an insufficient amount of storage to migrate data.

For more information

- [Remove failed drives from the cluster](#)
- [Basic MDSS drive troubleshooting](#)
- [Remove MDSS drives](#)
- [Replacing drives for SolidFire storage nodes](#)
- [Replacing drives for H600S series storage nodes](#)
- [H410S and H610S hardware information](#)

- [SF-series hardware information](#)

Remove failed drives from the cluster

The SolidFire system puts a drive in a failed state if the drive's self-diagnostics tells the node it has failed or if communication with the drive stops for five and a half minutes or longer. The system displays a list of the failed drives. You must remove a failed drive from the failed drive list in NetApp Element software.

Drives in the **Alerts** list show as **blockServiceUnhealthy** when a node is offline. When restarting the node, if the node and its drives come back online within five and a half minutes, the drives automatically update and continue as active drives in the cluster.

1. In the Element UI, select **Cluster > Drives**.
2. Click **Failed** to view the list of failed drives.
3. Note the slot number of the failed drive.

You need this information to locate the failed drive in the chassis.

4. Remove the failed drives using one of the following methods:

Option	Steps
To remove individual drives	a. Click Actions for the drive you want to remove. b. Click Remove.
To remove multiple drives	a. Select all the drives you want to remove, and click Bulk Actions. b. Click Remove.

Basic MDSS drive troubleshooting

You can recover metadata (or slice) drives by adding them back to the cluster in the event that one or both metadata drives fail. You can perform the recovery operation in the NetApp Element UI if the MDSS feature is already enabled on the node.

If either or both of the metadata drives in a node experiences a failure, the slice service will shut down and data from both drives will be backed up to different drives in the node.

The following scenarios outline possible failure scenarios, and provide basic recommendations to correct the issue:

System slice drive fails

- In this scenario, the slot 2 is verified and returned to an available state.
- The system slice drive must be repopulated before the slice service can be brought back online.
- You should replace the system slice drive, when the system slice drive becomes available, add the drive and the slot 2 drive at the same time.



You cannot add the drive in slot 2 by itself as a metadata drive. You must add both drives back to the node at the same time.

Slot 2 fails

- In this scenario, the system slice drive is verified and returned to an available state.
- You should replace slot 2 with a spare, when slot 2 becomes available, add the system slice drive and the slot 2 drive at the same time.

System slice drive and slot 2 fails

- You should replace both system slice drive and slot 2 with a spare drive. When both drives become available, add the system slice drive and the slot 2 drive at the same time.

Order of operations

- Replace the failed hardware drive with a spare drive (replace both drives if both have failed).
- Add drives back to the cluster when they have been repopulated and are in an available state.

Verify operations

- Verify that the drives in slot 0 (or internal) and slot 2 are identified as metadata drives in the Active Drives list.
- Verify that all slice balancing has completed (there are no further moving slices messages in the event log for at least 30 minutes).

For more information

[Add MDSS drives](#)

Add MDSS drives

You can add a second metadata drive on a SolidFire node by converting the block drive in slot 2 to a slice drive. This is accomplished by enabling the multi-drive slice service (MDSS) feature. To enable this feature, you must contact NetApp Support.

Getting a slice drive into an available state might require replacing a failed drive with a new or spare drive. You must add the system slice drive at the same time you add the drive for slot 2. If you try to add the slot 2 slice drive alone or before you add the system slice drive, the system will generate an error.

1. Click **Cluster > Drives**.
2. Click **Available** to view the list of available drives.
3. Select the slice drives to add.
4. Click **Bulk Actions**.
5. Click **Add**.
6. Confirm from the **Active Drives** tab that the drives have been added.

Remove MDSS drives

You can remove the multi-drive slice service (MDSS) drives. This procedure applies only if the node has multiple slice drives.



If the system slice drive and the slot 2 drive fail, the system will shutdown slice services and remove the drives. If there is no failure and you remove the drives, both drives must be removed at the same time.

1. Click **Cluster > Drives**.
2. From the **Available** drives tab, click the check box for the slice drives being removed.
3. Click **Bulk Actions**.
4. Click **Remove**.
5. Confirm the action.

Troubleshoot nodes

Remove nodes from a cluster

You can remove nodes from a cluster for maintenance or replacement. You should use the NetApp Element UI or API to remove nodes before taking them offline.

An overview of the procedure to remove storage nodes is as follows:

- Ensure that there is sufficient capacity in the cluster to create a copy of the data on the node.
- Remove drives from the cluster by using the UI or the RemoveDrives API method.

This results in the system migrating data from the node's drives to other drives in the cluster. The time this process takes is dependent on how much data must be migrated.

- Remove the node from the cluster.

Keep the following considerations in mind before you power down or power up a node:

- Powering down nodes and clusters involves risks if not performed properly.

Powering down a node should be done under the direction of NetApp Support.

- If a node has been down longer than 5.5 minutes under any type of shutdown condition, Double Helix data protection begins the task of writing single replicated blocks to another node to replicate the data. In this case, contact NetApp Support for help with analyzing the failed node.
- To safely reboot or power down a node, you can use the Shutdown API command.
- If a node is in a down, or in an off state, you must contact NetApp Support before bringing it back online.
- After a node is brought back online, you must add the drives back to the cluster, depending on how long it has been out of service.

For more information

[Replacing a failed SolidFire chassis](#)

[Replacing a failed H600S series node](#)

Power down a cluster

Perform the following procedure to power down an entire cluster.

Steps

1. (Optional) Contact NetApp Support for assistance with completing the preliminary steps.
2. Verify that all I/O has stopped.
3. Disconnect all iSCSI sessions:
 - a. Navigate to the management virtual IP (MVIP) address on the cluster to open the Element UI.
 - b. Note the nodes listed in the Nodes list.
 - c. Run the Shutdown API method with the halt option specified on each Node ID in the cluster.

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

[How to gracefully shut down and power on a NetApp Solidfire/HCI storage cluster](#)

Work with per-node utilities for storage nodes

Work with per-node utilities for storage nodes

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

Find more information

- [Access per-node settings using the per-node UI](#)
- [Network settings details from the per-node UI](#)
- [Cluster settings details from the per-node UI](#)
- [Run system tests using the per-node UI](#)

- [Run system utilities using the per-node UI](#)

Access per-node settings using the per-node UI

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

If you want to modify settings of a node in an Active state that is part of a cluster, you must log in as a cluster administrator user.



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

1. Open the per-node UI using one of the following methods:
 - Enter the management IP address followed by :442 in a browser window, and log in using an admin user name and password.
 - In the Element UI, select **Cluster > Nodes**, and click the management IP address link for the node you want to configure or modify. In the browser window that opens, you can edit the settings of the node.

NetApp

Hybrid Cloud Control

Node01

Node01

NETWORK SETTINGS

CLUSTER SETTINGS

SYSTEM TESTS

SYSTEM UTILITIES

Network Settings

Bond1G

Bond10G

Reset Changes

Method

static

Link Speed

1000

IPv4 Address

IPv4 Subnet Mask

255.255.255.0

IPv4 Gateway Address

IPv6 Address

IPv6 Gateway Address

MTU

1500

DNS Servers

Search Domains

Bond Mode

Status

Network settings details from the per-node UI

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

You can see the network settings for a storage node on the **Network Settings** page when you log in to the node (https://<node_IP>:442/hcc/node/network-settings). You can select either **Bond1G** (management) or **Bond10G** (storage) settings. The following list describes the settings that you can modify when a storage node is in Available, Pending, or Active state:

- **Method**

The method used to configure the interface. Possible methods:

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

- **Link Speed**

The speed negotiated by the virtual NIC.

- **IPv4 Address**

The IPv4 address for the eth0 network.

- **IPv4 Subnet Mask**

Address subdivisions of the IPv4 network.

- **IPv4 Gateway Address**

Router network address to send packets out of the local network.

- **IPv6 Address**

The IPv6 address for the eth0 network.

- **IPv6 Gateway Address**

Router network address to send packets out of the local network.

- **MTU**

Largest packet size that a network protocol can transmit. Must be greater than or equal to 1500. If you add a second storage NIC, the value should be 9000.

- **DNS Servers**

Network interface used for cluster communication.

- **Search Domains**

Search for additional MAC addresses available to the system.

- **Bond Mode**

Can be one of the following modes:

- ActivePassive (default)
- ALB
- LACP

- **Status**

Possible values:

- UpAndRunning
- Down
- Up

- **Virtual Network Tag**

Tag assigned when the virtual network was created.

- **Routes**

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

Cluster settings details from the per-node UI

You can verify cluster settings for a storage node after cluster configuration and modify the node hostname.

The following list describes the cluster settings for a storage node indicated from the **Cluster Settings** page of the per-node UI (https://<node_IP>:442/hcc/node/cluster-settings).

- **Role**

Role the node has in the cluster. Possible values:

- Storage: Storage or Fibre Channel node.
- Management: Node is a management node.

- **Hostname**

Name of the node.

- **Cluster**

Name of the cluster.

- **Cluster Membership**

State of the node. Possible values:

- Available: The node has no associated cluster name and is not yet part of a cluster.
- Pending: The node is configured and can be added to a designated cluster. Authentication is not required to access the node.
- PendingActive: The system is in the process of installing compatible software on the node. When complete, the node will move to the Active state.
- Active: The node is participating in a cluster. Authentication is required to modify the node.

- **Version**

Version of the Element software running on the node.

- **Ensemble**

Nodes that are part of the database ensemble.

- **Node ID**

ID assigned when a node is added to the cluster.

- **Cluster Interface**

Network interface used for cluster communication.

- **Management Interface**

Management network interface. This defaults to Bond1G but can also use Bond10G.

- **Storage Interface**

Storage network interface using Bond10G.

- **Encryption Capable**

Indicates whether or not the node supports drive encryption.

Run system tests using the per-node UI

You can test changes to the network settings after you commit them to the network configuration. You can run the tests to ensure that the storage node is stable and can be brought online without any issues.

You have logged in to the per-node UI for the storage node.

1. Click **System Tests**.
2. Click **Run Test** next to the test you want to run or select **Run All Tests**.



Running all test operations can be time consuming and should be done only at the direction of NetApp Support.

- **Test Connected Ensemble**

Tests and verifies the connectivity to a database ensemble. By default, the test uses the ensemble for the cluster the node is associated with. Alternatively you can provide a different ensemble to test connectivity.

- **Test Connect Mvip**

Pings the specified management virtual IP (MVIP) address and then executes a simple API call to the MVIP to verify connectivity. By default, the test uses the MVIP for the cluster the node is associated with.

- **Test Connect Svip**

Pings the specified storage virtual IP (SVIP) address using Internet Control Message Protocol (ICMP) packets that match the Maximum Transmission Unit (MTU) size set on the network adapter. It then connects to the SVIP as an iSCSI initiator. By default, the test uses the SVIP for the cluster the node is associated with.

- **Test Hardware Config**

Tests that all hardware configurations are correct, validates firmware versions are correct, and confirms all drives are installed and running properly. This is the same as factory testing.



This test is resource intensive and should only be run if requested by NetApp Support.

- **Test Local Connectivity**

Tests the connectivity to all of the other nodes in the cluster by pinging the cluster IP (CIP) on each node. This test will only be displayed on a node if the node is part of an active cluster.

- **Test Locate Cluster**

Validates that the node can locate the cluster specified in the cluster configuration.

- **Test Network Config**

Verifies that the configured network settings match the network settings being used on the system. This test is not intended to detect hardware failures when a node is actively participating in a cluster.

- **Test Ping**

Pings a specified list of hosts or, if none are specified, dynamically builds a list of all registered nodes in the cluster and pings each for simple connectivity.

- **Test Remote Connectivity**

Tests the connectivity to all nodes in remotely paired clusters by pinging the cluster IP (CIP) on each node. This test will only be displayed on a node if the node is part of an active cluster.

Run system utilities using the per-node UI

You can use the per-node UI for the storage node to create or delete support bundles, reset configuration settings for drives, and restart network or cluster services.

You have logged in to the per-node UI for the storage node.

1. Click **System Utilities**.
2. Click the button for the system utility that you want to run.

- **Control Power**

Reboots, power cycles, or shuts down the node.



This operation causes temporary loss of networking connectivity.

Specify the following parameters:

- Action: Options include Restart and Halt (power off).
- Wakeup Delay: Any additional time before the node comes back online.

- **Collect Node Logs**

Creates a support bundle under the node's /tmp/bundles directory.

Specify the following parameters:

- **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.
- **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.
- **Timeout Sec:** Specify the number of seconds to wait for each individual ping response.

◦ **Delete Node Logs**

Deletes any current support bundles on the node that were created using **Create Cluster Support Bundle** or the CreateSupportBundle API method.

◦ **Reset Drives**

Initializes drives and removes all data currently residing on the drive. You can reuse the drive in an existing node or in an upgraded node.

Specify the following parameter:

- **Drives:** List of device names (not driveIDs) to reset.

◦ **Reset Network Config**

Helps resolve network configuration issues for an individual node and resets an individual node's network configuration to the factory default settings.

◦ **Reset Node**

Resets a node to the factory settings. All data is removed but network settings for the node are preserved during this operation. Nodes can only be reset if they are unassigned to a cluster and in Available state.



All data, packages (software upgrades), configurations, and log files are deleted from the node when you use this option.

◦ **Restart Networking**

Restarts all networking services on a node.



This operation can cause temporary loss of network connectivity.

◦ **Restart Services**

Restarts Element software services on a node.



This operation can cause temporary node service interruption. You should perform this operation only at the direction of NetApp Support.

Specify the following parameters:

- **Service:** Service name to be restarted.

- Action: Action to perform on the service. Options include start, stop and restart.

Work with the management node

You can use the management node (mNode) to upgrade 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.



As a best practice, only associate one management node with one VMware vCenter instance, and avoid defining the same storage and compute resources or vCenter instances in multiple management nodes.

See [management node documentation](#) for more information.

Understand cluster fullness levels

The cluster running Element software generates cluster faults to warn the storage administrator when the cluster is running out of capacity. There are three levels of cluster fullness, all of which are displayed in the NetApp Element UI: warning, error, and critical.

The system uses the BlockClusterFull error code to warn about cluster block storage fullness. You can view the cluster fullness severity levels from the Alerts tab of the Element UI.

The following list includes information about the BlockClusterFull severity levels:

- **Warning**

This is a customer-configurable warning that appears when the cluster's block capacity is approaching the error severity level. By default, this level is set at three percent under the error level and can be tuned via the Element UI and API. You must add more capacity, or free up capacity as soon as possible.

- **Error**

When the cluster is in this state, if a node is lost, there will not be enough capacity in the cluster to rebuild Double Helix data protection. New volume creation, clones, and snapshots are all blocked while the cluster is in this state. This is not a safe or recommended state for any cluster to be in. You must add more capacity, or free up capacity immediately.

- **Critical**

This critical error has occurred because the cluster is 100 percent consumed. It is in a read-only state and no new iSCSI connections can be made to the cluster. When this stage is reached, you must free up or add more capacity immediately.

The system uses the MetadataClusterFull error code to warn about cluster metadata storage fullness. You can view the cluster metadata storage fullness from the Cluster Capacity section on the Overview page of the Reporting tab in the Element UI.

The following list includes information about the MetadataClusterFull severity levels:

- **Warning**

This is a customer-configurable warning that appears when the cluster's metadata capacity is approaching the error severity level. By default, this level is set at three percent under the error level and can be tuned via the Element API. You must add more capacity, or free up capacity as soon as possible.

- **Error**

When the cluster is in this state, if a node is lost, there will not be enough capacity in the cluster to rebuild Double Helix data protection. New volume creation, clones, and snapshots are all blocked while the cluster is in this state. This is not a safe or recommended state for any cluster to be in. You must add more capacity, or free up capacity immediately.

- **Critical**

This critical error has occurred because the cluster is 100 percent consumed. It is in a read-only state and no new iSCSI connections can be made to the cluster. When this stage is reached, you must free up or add more capacity immediately.



The following applies to two-node cluster thresholds:

- Metadata fullness error is 20% below critical.
- Block fullness error is 1 block drive (including stranded capacity) below critical; meaning that it is two block drives worth of capacity below critical.

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—without prior written permission of the copyright owner.

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

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

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

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

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.