



Monitor system performance

VCP

NetApp
November 18, 2025

This PDF was generated from https://docs.netapp.com/us-en/vcp/vcp_task_reports_intro.html on November 18, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Monitor system performance	1
Monitor system performance with Reporting options	1
Find more information	1
Monitor overall cluster health on the Overview page	1
Reporting Overview page data	1
Find more information	3
Monitor system alerts	3
Find more information	4
System alerts list	4
Monitor event logs for troubleshooting	20
Event types	21
Find more information	22
Monitor volume performance	22
Volume performance data	23
Find more information	23
Monitor iSCSI sessions to determine connection status	23
iSCSI session data	24
Find more information	24
Monitor VM performance tiering with QoSSIOC events	24
QoSSIOC event data	25
Find more information	25

Monitor system performance

Monitor system performance with Reporting options

You can view information about the cluster's components and performance by using the Reporting pages of the NetApp Element Plug-in for VMware vCenter Server.

Using the vCenter Plug-in, you can monitor cluster components and performance in the following ways:

- [Monitor overall cluster health on the Overview page](#)
- [Monitor system alerts](#)
- [Monitor event logs for troubleshooting](#)
- [Monitor volume performance](#)
- [Monitor iSCSI sessions to determine connection status](#)
- [Monitor VM performance tiering with QoSIOC events](#)

Find more information

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

Monitor overall cluster health on the Overview page

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

Steps

1. From the vCenter Plug-in, open the **Reporting** tab:
 - Beginning with Element vCenter plug-in 5.0, select **NetApp Element Remote Plugin > Management > Reporting**.
 - For Element vCenter plug-in 4.10 and earlier, select **NetApp Element Management > Reporting**.
2. Monitor the data on the **Overview** page.

Reporting Overview page data

The following data appears on the Reporting Overview page:

- **Cluster Capacity:** The capacity remaining for block storage, metadata, and provisioned space. Move the pointer over the progress bar to see threshold information.
- **Cluster Information:** Information specific to the cluster, such as cluster name, the version of NetApp Element software running on the cluster, MVIP and SVIP addresses, and the number of nodes, 4k IOPS, volumes, and sessions on the cluster.
 - **Cluster Name:** The name for the cluster.

- **Storage IP (SVIP):** The storage virtual IP address (SVIP).
- **Management IP (MVIP):** The management virtual IP address (MVIP).
- **SVIP VLAN Tag:** The VLAN identifier for the master SVIP address.
- **MVIP VLAN Tag:** The VLAN identifier for the master MVIP address.
- **Node Count:** The number of active nodes in the cluster.
- **Cluster 4K IOPS:** The number of 4096 (4K) blocks that can be read/written by the cluster in a second.
- **Element OS Version:** The version of the NetApp Element software that the cluster is running.
- **Volume Count:** The total number of volumes, excluding virtual volumes, on the cluster.
- **Virtual Volume Count:** The total number of virtual volumes on the cluster.
- **iSCSI Sessions:** The iSCSI sessions that are connected to the cluster.
- **Fibre Channel Sessions:** The Fibre Channel sessions that are connected to the cluster.
- **Cluster Efficiency:** Overall system capacity that is being utilized that takes into account thin provisioning, deduplication, and compression. The calculated benefit achieved on the cluster is calculated by comparing what the capacity utilization would be without thin provisioning, deduplication, and compression on a traditional storage device.
- **Protection Domains:** A summary of protection domains monitoring for the cluster.



The protection domains feature is not compatible with two-node clusters.

- **Protection Domains Monitoring Level:** The protection domain resiliency levels as selected by the user. Possible values are Chassis or Node. Green indicates that the cluster is capable of the selected monitoring level. Red indicates that the cluster is no longer capable of the selected monitoring level and corrective action is needed.
- **Remaining Block Capacity:** Indicates the percentage of block capacity that is remaining to maintain the selected resiliency level.
- **Metadata Capacity:** Indicates if there is sufficient metadata capacity to heal from failure while also maintaining undisrupted data availability. Normal (green) indicates that the cluster has sufficient metadata to maintain the selected monitoring level. Full (red) indicates that the cluster is no longer capable of the selected monitoring level and corrective action is needed.
- **Custom Protection Domain Health:** Displays the custom Protection Domain health status for the cluster when a custom Protection Domain is configured on the cluster.

The following data indicates the protection available against the failure of one of the custom Protection Domains for the cluster.

- **Protection Level:** Indicates the overall protection level status.
- **Block Capacity:** Indicates the current protection level status of the block services subsystem.

It also indicates the total capacity threshold at which resiliency is lost.

- **Metadata Capacity:** Indicates the current protection level status of the metadata services subsystem.
- **Ensemble Nodes:** Indicates the current protection level status of the ensemble members subsystem.

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



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

Minimum IOPS: $2000/50000 = 0.04x$

Maximum IOPS: $60000/50000 = 1.20x$

Burst IOPS: $60000/50000 = 1.20x 1.00x$

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

- **Cluster Health:** The hardware, capacity, and security components of the health of the cluster. Color codes indicate the following:
 - **Green:** Healthy
 - **Yellow:** Critical
 - **Red:** Error
- **Cluster Input/Output:** The I/O currently running on the cluster. The values are calculated based on the previous I/O measurement against the current I/O measurements. These are the measurements shown in the graph:
 - **Total:** The combined read and write IOPS occurring in the system.
 - **Read:** The number of read IOPS occurring.
 - **Write:** The number of write IOPS.
- **Cluster Throughput:** The bandwidth activity for read, write, and total bandwidth on the cluster:
 - **Total:** The total MB/s used for both read and write activity in the cluster.
 - **Read:** The read activity in MB/s for the cluster.
 - **Write:** The write activity in MB/s for the cluster.
- **Performance Utilization:** The percentage of cluster IOPS being consumed. For example, a 250K IOPS cluster running at 100K IOPS would show 40% consumption.

Find more information

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

Monitor system alerts

You can monitor alerts, which are information, warnings, or errors that indicate how well the cluster is running.

Alerts are cluster faults or errors and are reported as they occur. Most errors resolve themselves automatically; however, some might require manual intervention. The system reports alert error codes with each alert on the Alerts page. Error codes help you determine what component of the system experienced the alert and why the alert was generated. See [System alerts list](#) for a descriptions and remediation steps.

After you resolve the issue, the system polls itself and identifies the issue as resolved. Then, all information about the alert including the date it was resolved is moved to the Resolved view.

Steps

1. From the vCenter Plug-in, open the **Management** tab:
 - Beginning with Element vCenter plug-in 5.0, select **NetApp Element Remote Plugin > Management > Management**.
 - For Element vCenter plug-in 4.10 and earlier, select **NetApp Element Management > Management**.
2. Select **Reporting > Alerts**.
3. Monitor the following cluster alert information:
 - **ID**: Unique ID for a cluster alert.
 - **Severity**
 - **warning**: A minor issue that might soon require attention. System upgrades are still allowed at this severity level.
 - **error**: A failure that might cause performance degradation or loss of high availability (HA). Errors generally should not affect service otherwise.
 - **critical**: A serious failure that affects service. The system is unable to serve API or client I/O requests. Operating in this state could lead to potential loss of data.
 - **bestPractice**: A recommended system configuration best practice is not being used.
 - **Type**
 - **node**: Fault affecting an entire node.
 - **drive**: Fault affecting an individual drive.
 - **cluster**: Fault affecting the entire cluster.
 - **service**: Fault affecting a service on the cluster.
 - **volume**: Fault affecting a volume on the cluster.
 - **Node**: Node ID for the node that this fault refers to. Included for node and drive faults, otherwise set to - (dash).
 - **Drive ID**: Drive ID for the drive that this fault refers to. Included for drive faults, otherwise set to - (dash).
 - **Error Code**: A descriptive code that indicates what caused the fault.
 - **Details**: Detailed description of the fault.
 - **Time**: This heading is visible only in the Active filter view. The date and time the fault was logged.
 - **Resolution Date**: This heading is visible only in the Resolved filter view. The date and time the fault was resolved.
4. To validate that the issue was resolved, look for it in the Resolved view.

Find more information

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

System alerts list

The system reports error codes with each alert that help you determine what component of the system experienced the alert and why the alert was generated. You can view the error codes using the plug-in extension point:

- Beginning with Element vCenter plug-in 5.0, select **NetApp Remote Plugin > Management > Reporting > Alerts**.
- For Element vCenter plug-in 4.10 and earlier, select **NetApp Element Management > Reporting > Alerts**.

The following list outlines the different types of system alerts.

- **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): 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 `TestKeyServerKmip` 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 cErrorCount reaches cDimmCorrectableErrWarnThreshold.	Warning	Correctable ECC memory errors above threshold on DIMM: <Processor> <DIMM Slot>
A single DIMM cErrorCount stays above cDimmCorrectableErrWarnThreshold until cErrorFaultTimer expires for the DIMM.	Error	Correctable ECC memory errors above threshold on DIMM: <Processor> <DIMM>
A memory controller reports cErrorCount above cMemCtrlrCorrectableErrWarnThreshold, and cMemCtrlrCorrectableErrWarnDuration is specified.	Warning	Correctable ECC memory errors above threshold on memory controller: <Processor> <Memory Controller>

A memory controller reports cErrorCount above cMemCtrlrCorrectableErrWarnThreshold 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 cMemCtrlrUncorrectableErrFaultThreshold.	Warning	Uncorrectable ECC memory error(s) detected on memory controller: <Processor> <Memory Controller>
A memory controller reports a uErrorCount of at least cMemCtrlrUncorrectableErrFaultThreshold.	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.

- **nvrampDeviceStatus**

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

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

Warning	<p>A warning has been detected by the hardware. This condition may be transitory, such as a temperature warning.</p> <ul style="list-style-type: none"> • nvmLifetimeError • nvmLifetimeStatus • energySourceLifetimeStatus • energySourceTemperatureStatus • warningThresholdExceeded
Error	<p>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:</p> <ul style="list-style-type: none"> • NVRAM device mount point doesn't exist. • NVRAM device partition doesn't exist. • NVRAM device partition exists, but not mounted.
Critical	<p>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.</p> <ul style="list-style-type: none"> • 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.

Monitor event logs for troubleshooting

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

Steps

1. From the vCenter Plug-in, open the **Management** tab:
 - Beginning with Element vCenter plug-in 5.0, select **NetApp Element Remote Plugin > Management > Management**.
 - For Element vCenter plug-in 4.10 and earlier, select **NetApp Element Management > Management**.
2. Select **Reporting > Event Log**.
3. To review details, select an event and click **Details**.
4. Review the event information that includes the following:
 - **Event Type**: The type of event being logged; for example, API events or clone events.
 - **Service ID**: The ID of the service that reported the event (if applicable). The value is zero if the fault is not associated with a service.
 - **Node or Drive ID**: The ID of the node or drive that reported the event (if applicable).

Event types

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



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

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

- **apiEvent**: Events initiated by a user through an API or web UI that modify settings.
- **binAssignmentsEvent**: Events related to the assignment of data bins. Bins are essentially containers that hold data and are mapped across the cluster.
- **binSyncEvent**: System events related to a reassignment of data among block services.
- **bsCheckEvent**: System events related to block service checks.
- **bsKillEvent**: System events related to block service terminations.
- **bulkOpEvent**: Events related to operations performed on an entire volume, such as a backup, restore, snapshot, or clone.
- **cloneEvent**: Events related to volume cloning.
- **clusterMasterEvent**: Events appearing upon cluster initialization or upon configuration changes to the cluster, such as adding or removing nodes.
- **csumEvent**: Events related to invalid data checksums on the disk.
- **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 status of virtual networking.
- **platformHardwareEvent**: Events related to issues detected on hardware devices.
- **remoteClusterEvent**: Events related to remote cluster pairing.
- **schedulerEvent**: Events related to scheduled snapshots.
- **serviceEvent**: Events related to system service status.

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

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.

Find more information

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

Monitor volume performance

You can view performance information for all volumes in the selected cluster from Reporting tab of the plug-in extension point.

Steps

1. From the vCenter Plug-in, open the **Reporting** tab:
 - Beginning with Element vCenter plug-in 5.0, select **NetApp Element Remote Plugin > Management > Reporting**.
 - For Element vCenter plug-in 4.10 and earlier, select **NetApp Element Management > Reporting**.
2. Select **Volume Performance**.
3. To change how often the data refreshes on the page, click **Refresh every list** and choose a 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.

Volume performance data

- **Name:** Name of the volume when it was created.
- **Account:** The name of the account assigned to the volume.
- **Access Groups:** The name of the volume access group or groups to which the volume belongs.
- **Volume Utilization %:** A percentage value that describes how much the client is using the volume.

Possible values:

- 0 = Client is not using the volume
- 100 = Client is using the max
- >100 = Client is using the burst
- **Total IOPS:** The total number of IOPS (read and write) currently being executed against the volume.
- **Read IOPS:** The total number of read IOPS currently being executed against the volume.
- **Write IOPS:** The total number of write IOPS currently being executed against the volume.
- **Total Throughput:** The total amount of throughput (read and write) currently being executed against the volume.
- **Read Throughput:** The total amount of read throughput currently being executed against the volume.
- **Write Throughput:** The total amount of write throughput currently being executed against the volume.
- **Total Latency (ms):** The average time, in microseconds, to complete read and write operations to a volume.
- **Read Latency (ms):** The average time, in microseconds, to complete read operations to the volume in the last 500 milliseconds.
- **Write Latency (ms):** The average time, in microseconds, to complete write operations to a volume in the last 500 milliseconds.
- **Queue Depth:** The number of outstanding read and write operations to the volume.
- **Average IO Size:** Average size in bytes of recent I/O to the volume in the last 500 milliseconds.

Find more information

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

Monitor iSCSI sessions to determine connection status

You can view information about iSCSI sessions that are connected to the selected cluster in the NetApp Element Plug-in for VMware vCenter Server.

Steps

1. From the vCenter Plug-in, open the **Reporting** tab:
 - Beginning with Element vCenter plug-in 5.0, select **NetApp Element Remote Plugin > Management > Reporting**.
 - For Element vCenter plug-in 4.10 and earlier, select **NetApp Element Management > Reporting**.

2. Select **iSCSI Sessions**.

iSCSI session data

- **Node**: The node hosting the primary metadata partition for the volume.
- **Account**: The name of the account that owns the volume. If value is blank, a dash (-) will be displayed.
- **Volume**: The volume name identified on the node.
- **Volume ID**: ID of the volume associated with the Target IQN.
- **Initiator ID**: A system-generated ID for the initiator.
- **Initiator Alias**: An optional name for the initiator that makes finding the initiator easier in a long list.
- **Initiator IP**: The IP address of the endpoint that initiates the session.
- **Initiator IQN**: The IQN of the endpoint that initiates the session.
- **Target IP**: The IP address of the node hosting the volume.
- **Target IQN**: The IQN of the volume.
- **Created On**: Date the session was established.

Find more information

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

Monitor VM performance tiering with QoSSIOC events

You can view events related to QoSSIOC when a VM with a QoS-enabled datastore is reconfigured or issued a power or guest event.

You can view QoSSIOC events from the plug-in extension point in NetApp Element Plug-in for vCenter Server.

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

 (i)

- Beginning with Element vCenter plug-in 5.0, to use [vCenter Linked Mode](#), you register the Element Plug-in from a separate management node for each vCenter Server that manages NetApp SolidFire storage clusters.
- Using NetApp Element Plug-in for vCenter Server 4.10 and earlier to manage cluster resources from other vCenter Servers using [vCenter Linked Mode](#) is limited to local storage clusters only.

What you'll need

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

Steps

1. In your vSphere Web Client, open the **QoSSIOC Events** tab:
 - Beginning with Element vCenter plug-in 5.0, select **NetApp Element Remote Plugin > Configuration > QoSSIOC Events**.
 - For Element vCenter plug-in 4.10 and earlier, select **NetApp Element Configuration > QoSSIOC Events**.

QoSSIOC event data

- **Date**: The date and time of the QoSSIOC event.
- **Datastore Name**: The user-defined datastore name.
- **Cluster IP**: The IP address of the cluster containing the datastore from which the event originated.
- **Volume ID**: The system-generated ID for the associated volume.
- **Min IOPs**: The current minimum IOPS QoS setting of the volume.
- **Max IOPs**: The current maximum IOPS QoS setting of the volume.
- **Burst IOPs**: The current maximum burst QoS setting of the volume.
- **Burst Time**: The length of time a burst is allowed.

Find more information

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

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.