



Monitor and manage cluster health

Active IQ Unified Manager 9.13

NetApp
October 16, 2025

Table of Contents

- Monitor and manage cluster health 1
 - Introduction to Active IQ Unified Manager health monitoring 1
 - Physical and logical capacity 1
 - Capacity measurement units 1
 - Unified Manager health monitoring features 2
 - Unified Manager interfaces used to manage storage system health 3
- Managing and monitoring clusters and cluster object health 4
 - Understanding cluster monitoring 4
 - Viewing the cluster list and details 5
 - Checking the health of clusters in a MetroCluster configuration 6
 - Viewing the health and capacity status of All SAN Array clusters 8
 - Viewing the node list and details 8
 - Generating a hardware inventory report for contract renewal 9
 - Viewing the Storage VM list and details 9
 - Viewing the aggregate list and details 10
 - Viewing FabricPool capacity information 11
 - Viewing storage pool details 12
 - Viewing the volume list and details 13
 - Viewing details about NFS shares 13
 - Viewing details about SMB/CIFS shares 14
 - Viewing the list of Snapshot copies 15
 - Deleting Snapshot copies 15
 - Calculating reclaimable space for Snapshot copies 16
 - Description of cluster object windows and dialog boxes 17
- Common Unified Manager health workflows and tasks 17
 - Monitoring and troubleshooting data availability 18
 - Resolving capacity issues 25
 - Managing health thresholds 26
 - Managing cluster security objectives 31
 - Managing backup and restore operations 43
 - Managing scripts 58
 - Managing and monitoring groups 61
 - Prioritizing storage object events using annotations 70
 - Sending a support bundle through web UI and maintenance console 78
 - Tasks and information related to several workflows 85

Monitor and manage cluster health

Introduction to Active IQ Unified Manager health monitoring

Active IQ Unified Manager (formerly OnCommand Unified Manager) helps you to monitor a large number of systems running ONTAP software through a centralized user interface. The Unified Manager server infrastructure delivers scalability, supportability, and enhanced monitoring and notification capabilities.

The key capabilities of Unified Manager include monitoring, alerting, managing availability and capacity of clusters, managing protection capabilities, and bundling of diagnostic data and sending it to technical support.

You can use Unified Manager to monitor your clusters. When issues occur in the cluster, Unified Manager notifies you about the details of such issues through events. Some events also provide you with a remedial action that you can take to rectify the issues. You can configure alerts for events so that when issues occur, you are notified through email, and SNMP traps.

You can use Unified Manager to manage storage objects in your environment by associating them with annotations. You can create custom annotations and dynamically associate clusters, storage virtual machines (SVMs), and volumes with the annotations through rules.

You can also plan the storage requirements of your cluster objects using the information provided in the capacity and health charts, for the respective cluster object.

Physical and logical capacity

Unified Manager makes use of the concepts of physical and logical space used for ONTAP storage objects.

- **Physical capacity:** Physical space refers to the physical blocks of storage used in the volume. “Physical used capacity” is typically smaller than logical used capacity due to the reduction of data from storage efficiency features (such as deduplication and compression).
- **Logical capacity:** Logical space refers to the usable space (the logical blocks) in a volume. Logical space refers to how theoretical space can be used, without accounting for results of deduplication or compression. “Logical space used” is physical space used plus the savings from storage efficiency features (such as deduplication and compression) that have been configured. This measurement often appears larger than the physical used capacity because this does not reflect the data compression and other reductions in the physical space. Thus, the total logical capacity could be higher than the provisioned space.

Capacity measurement units

Unified Manager calculates storage capacity based on binary units of 1024 (2^{10}) bytes. In ONTAP 9.10.0 and earlier, these units were displayed as KB, MB, GB, TB, and PB. Beginning with ONTAP 9.10.1, they are displayed in Unified Manager as KiB, MiB, GiB, TiB, and PiB.



The units used for throughput continue to be kilobytes per second (Kbps), Megabytes per second (Mbps), Gigabytes per second (Gbps), or Terabytes per second (Tbps) and so forth, for all releases of ONTAP.

Capacity unit displayed in Unified Manager for ONTAP 9.10.0 and earlier	Capacity unit displayed in Unified Manager for ONTAP 9.10.1	Calculation	Value in bytes
KB	KiB	1024	1024 bytes
MB	MiB	1024 * 1024	1,048,576 bytes
GB	GiB	1024 * 1024 * 1024	1,073,741,824 bytes
TB	TiB	1024 * 1024 * 1024 * 1024	1,099,511,627,776 bytes

Unified Manager health monitoring features

Unified Manager is built on a server infrastructure that delivers scalability, supportability, and enhanced monitoring and notification capabilities. Unified Manager supports monitoring of systems running ONTAP software.

Unified Manager includes the following features:

- Discovery, monitoring, and notifications for systems that are installed with ONTAP software:
 - Physical objects: nodes, disks, disk shelves, SFO pairs, ports, and Flash Cache
 - Logical objects: clusters, storage virtual machines (SVMs), aggregates, volumes, LUNs, namespaces, qtrees, LIFs, Snapshot copies, junction paths, NFS shares, SMB shares, user and group quotas, QoS policy groups, and initiator groups
 - Protocols: CIFS, NFS, FC, iSCSI, NVMe, and FCoE
 - Storage efficiency: SSD aggregates, Flash Pool aggregates, FabricPool aggregates, deduplication, and compression
 - Protection: SnapMirror relationships (synchronous and asynchronous) and SnapVault relationships
- Viewing the cluster discovery and monitoring status
- MetroCluster over FC and IP configurations: viewing and monitoring the configuration, issues, and connectivity status of the cluster components. MetroCluster switches and bridges for MetroCluster over FC configurations
- Enhanced alerts, events, and threshold infrastructure
- LDAP, LDAPS, SAML authentication, and local user support
- RBAC (for a predefined set of roles)
- AutoSupport and support bundle
- Enhanced dashboard to show capacity, availability, protection, and performance health of the environment
- Volume move interoperability, volume move history, and junction path change history
- Scope of Impact area that graphically displays the resources that are impacted for events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events

- Possible Effect area that displays the effect of the MetroCluster events
- Suggested Corrective Actions area that displays the actions that can be performed to address events such as Some Failed Disks, MetroCluster Aggregate Mirroring Degraded, and MetroCluster Spare Disks Left Behind events
- Resources that Might be Impacted area that displays the resources that might be impacted for events such as for the Volume Offline event, the Volume Restricted event, and the Thin-Provisioned Volume Space At Risk event
- Support for SVMs with FlexVol or FlexGroup volumes
- Support for monitoring node root volumes
- Enhanced Snapshot copy monitoring, including computing reclaimable space and deleting Snapshot copies
- Annotations for storage objects
- Report creation and management of storage object information such as physical and logical capacity, utilization, space savings, performance, and related events
- Integration with OnCommand Workflow Automation to execute workflows

The Storage Automation Store contains NetApp-certified automated storage workflow packs developed for use with OnCommand Workflow Automation (WFA). You can download the packs, and then import them to WFA to execute them. The automated workflows are available here:

[Storage Automation Store](#)

Unified Manager interfaces used to manage storage system health

These sections contain information about the two user interfaces that Active IQ Unified Manager provides for troubleshooting data storage capacity, availability, and protection issues. The two UIs are the Unified Manager web UI and the maintenance console.

If you want to use the protection features in Unified Manager, you must also install and configure OnCommand Workflow Automation (WFA).

Unified Manager web UI

The Unified Manager web UI enables an administrator to monitor and troubleshoot cluster issues relating to data storage capacity, availability, and protection.

These sections describe some common workflows that an administrator can follow to troubleshoot storage capacity, data availability, or protection issues displayed in the Unified Manager web UI.

Maintenance console

The Unified Manager maintenance console enables an administrator to monitor, diagnose, and address operating system issues, version upgrade issues, user access issues, and network issues related to the Unified Manager server itself. If the Unified Manager web UI is unavailable, the maintenance console is the only form of access to Unified Manager.

You can use this information for accessing the maintenance console and using it to resolve issues related to the functioning of the Unified Manager server.

Managing and monitoring clusters and cluster object health

Unified Manager uses periodic API queries and a data collection engine to collect data from the clusters. By adding clusters to the Unified Manager database, you can monitor and manage these clusters for any availability and capacity risks.

Understanding cluster monitoring

You can add clusters to the Unified Manager database to monitor clusters for availability, capacity, and other details, such as CPU usage, interface statistics, free disk space, qtree usage, and chassis environmental.

Events are generated if the status is abnormal or when a predefined threshold is breached. If configured to do so, Unified Manager sends a notification to a specified recipient when an event triggers an alert.

Understanding node root volumes

You can monitor the node root volume using Unified Manager. The best practice is that the node root volume should have sufficient capacity to prevent the node from going down.

When the used capacity of the node root volume exceeds 80 percent of the total node root volume capacity, the Node Root Volume Space Nearly Full event is generated. You can configure an alert for the event to get a notification. You can take appropriate actions to prevent the node from going down by using either ONTAP System Manager or the ONTAP CLI.

Understanding events and thresholds for node root aggregates

You can monitor the node root aggregate by using Unified Manager. The best practice is to thickly provision the root volume in the root aggregate to prevent the node from halting.

By default, capacity and performance events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to the node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are modified by the technical support representative, the capacity threshold values are applied to the node root aggregate.

You can take appropriate actions to prevent the node from halting by using either ONTAP System Manager or the ONTAP CLI.

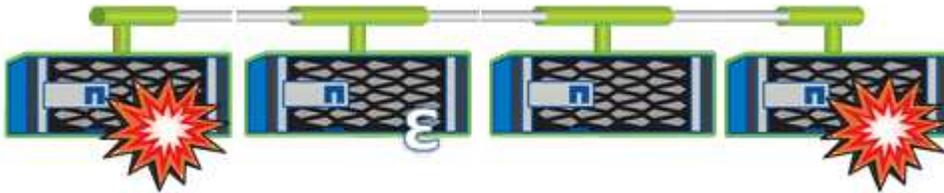
Understanding quorum and epsilon

Quorum and epsilon are important measures of cluster health and function that together indicate how clusters address potential communications and connectivity challenges.

Quorum is a precondition for a fully functioning cluster. When a cluster is in quorum, a simple majority of nodes are healthy and can communicate with each other. When quorum is lost, the cluster loses the ability to accomplish normal cluster operations. Only one collection of nodes can have quorum at any one time because all of the nodes collectively share a single view of the data. Therefore, if two non-communicating nodes are permitted to modify the data in divergent ways, it is no longer possible to reconcile the data into a single data view.

Each node in the cluster participates in a voting protocol that elects one node master; each remaining node is a secondary. The master node is responsible for synchronizing information across the cluster. When quorum is formed, it is maintained by continual voting. If the master node goes offline and the cluster is still in quorum, a new master is elected by the nodes that remain online.

Because there is the possibility of a tie in a cluster that has an even number of nodes, one node has an extra fractional voting weight called epsilon. If the connectivity between two equal portions of a large cluster fails, the group of nodes containing epsilon maintains quorum, assuming that all of the nodes are healthy. For example, the following illustration shows a four-node cluster in which two of the nodes have failed. However, because one of the surviving nodes holds epsilon, the cluster remains in quorum even though there is not a simple majority of healthy nodes.



Epsilon is automatically assigned to the first node when the cluster is created. If the node that holds epsilon becomes unhealthy, takes over its high-availability partner, or is taken over by its high-availability partner, then epsilon is automatically reassigned to a healthy node in a different HA pair.

Taking a node offline can affect the ability of the cluster to remain in quorum. Therefore, ONTAP issues a warning message if you attempt an operation that will either take the cluster out of quorum or else put it one outage away from a loss of quorum. You can disable the quorum warning messages by using the cluster quorum-service options modify command at the advanced privilege level.

In general, assuming reliable connectivity among the nodes of the cluster, a larger cluster is more stable than a smaller cluster. The quorum requirement of a simple majority of half the nodes plus epsilon is easier to maintain in a cluster of 24 nodes than in a cluster of two nodes.

A two-node cluster presents some unique challenges for maintaining quorum. Two-node clusters use cluster HA, in which neither node holds epsilon; instead, both nodes are continuously polled to ensure that if one node fails, the other has full read-write access to data, as well as access to logical interfaces and management functions.

Viewing the cluster list and details

You can use the Health: All Clusters view to view your inventory of clusters. The Capacity: All Clusters view enables you to view summarized information about storage capacity and utilization in all clusters.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

You can also view details for individual clusters such as the cluster health, capacity, configuration, LIFs, nodes, and disks in that cluster by using the Cluster / Health details page.

The details in the Health: All Clusters view, Capacity: All Clusters view, and the Cluster / Health details page help you plan your storage. For example, before provisioning a new aggregate, you can select a specific cluster from the Health: All Clusters view and obtain capacity details to determine if the cluster has the required space.

Steps

1. In the left navigation pane, click **Storage > Clusters**.
2. In the View menu, select **Health: All Clusters** view to view health information, or **Capacity: All Clusters** view to view details about storage capacity and utilization in all clusters.
3. Click the name of a cluster to view complete details of the cluster in the **Cluster / Health** details page.

Related information

- [Cluster / Health details page](#)
- [Performance: All Clusters view](#)
- [Monitoring MetroCluster configurations](#)
- [Viewing security status for clusters and Storage VMs](#)
- [What security criteria are being evaluated](#)

Checking the health of clusters in a MetroCluster configuration

You can use Active IQ Unified Manager (Unified Manager) to check the operational health of clusters, and their components, in MetroCluster over FC and MetroCluster over IP configurations. If the clusters were involved in a performance event detected by Unified Manager, the health status can help you determine whether a hardware or software issue contributed to the event.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have analyzed a performance event for a MetroCluster configuration and obtained the name of the cluster involved.
- Both clusters in the MetroCluster configuration over FC and IP must be monitored by the same instance of Unified Manager.

Determining cluster health in MetroCluster over FC configuration

Follow these steps for determining cluster health in a MetroCluster over FC configuration.

Steps


1. In the left navigation pane, click **Event Management** to display the event list.
2. In the filter panel, select all MetroCluster filters under the **Source Type** category. You see all events raised in your environment for all MetroCluster configurations.
3. Next to a MetroCluster event, click the name of the cluster.



If no MetroCluster events are displayed, you can use the Search bar to search for the name of the cluster involved in the event related to your MetroCluster over FC configuration.

The Health: All Clusters view is displayed with detailed information about the event.


4. Select the **MetroCluster Connectivity** tab to display the health of the connection between the selected cluster and its partner cluster.



In this example, the names and the components of the local cluster and its partner cluster are displayed. A yellow or red icon indicates a health event for the highlighted component. The Connectivity icon represents the link between the clusters. You can point your mouse cursor to an icon to display event information or click the icon to display the events. A health issue on either cluster might have contributed to the performance event.

Unified Manager monitors the NVRAM component of the link between the clusters. If the FC Switches icon on the local or partner cluster or the Connectivity icon is red, a link health issue might have caused the performance event.

5. Select the **MetroCluster Replication** tab.



In this example, if the NVRAM icon on the local or partner cluster is yellow or red, a health issue with the NVRAM might have caused the performance event. If there are no red or yellow icons on the page, a performance issue on the partner cluster might have caused the performance event.

Determining cluster health in MetroCluster over IP configuration

Follow these steps for determining cluster health in a MetroCluster over IP configuration.

Steps

1. In the left navigation pane, click **Event Management** to display the event list.
2. In the filter panel, under the **Source Type** category, select the `MetroCluster Relationship` filter. You see all events raised in your environment for all MetroCluster configurations.



If you cannot see the reported MetroCluster events, you can use the Search bar to search by the name of the cluster involved in the event related to your MetroCluster over IP configuration.

3. Next to the relevant MetroCluster event, click the name of the cluster. The Clusters page is displayed with the details of that cluster. For information about determining health issues, see [Monitor connectivity issues in MetroCluster over IP configuration](#).

Viewing the health and capacity status of All SAN Array clusters

You can use the Cluster inventory pages to display the health and capacity status of your All SAN Array clusters.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view overview information for All SAN Array clusters in the Health: All Clusters view and Capacity: All Clusters view. Additionally, you can view details in the Cluster / Health details page.

Steps

1. In the left navigation pane, click **Storage > Clusters**.
2. Make sure that the "Personality" column is displayed in the **Health: All Clusters** view, or add it using the **Show / Hide** control.

This column displays "All SAN Array" for your All SAN Array clusters.

3. Review the information.
4. To view information about storage capacity in those clusters, select the Capacity: All Clusters view.
5. To view detailed information about health and storage capacity in those clusters, click the name of an All SAN Array cluster.

View the details in the Health, Capacity, and Nodes tabs in the Cluster / Health details page

Viewing the node list and details

You can use the Health: All Nodes view to view the list of nodes in your clusters. You can use the Cluster / Health details page to view detailed information about nodes that are part of the cluster that is monitored.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view details such as the node state, cluster that contains the node, aggregate capacity details (used and total), and raw capacity details (usable, spare, and total). You can also obtain information about HA pairs, disks shelves, and ports.

Steps

1. In the left navigation pane, click **Storage > Nodes**.

2. On the **Health: All Nodes** view, click the node whose details you want to view.

The detailed information for the selected node is displayed in the Cluster / Health details page. The left pane displays the list of HA pairs. By default, the HA Details is open, which displays HA state details and events related to the selected HA pair.

3. To view other details about the node, perform the appropriate action:

To view...	Click...
Details about the disk shelves	Disk Shelves.
Port-related information	Ports.

For more information, see:

- [Performance: All Nodes view](#)
- [Viewing node and aggregate available IOPS values](#)
- [Viewing node and aggregate performance capacity used values](#)

Generating a hardware inventory report for contract renewal

You can generate a report that contains a complete list of cluster and node information; such as hardware model numbers and serial numbers, disk types and counts, installed licenses, and more. This report is helpful for contract renewal within secure sites (“dark” sites) that are not connected to the NetAppActive IQ platform.

What you’ll need

You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage > Nodes**.
2. Go to the **Health: All Nodes** view or **Performance: All Nodes** view.
3. Select **Reports > * > Hardware Inventory Report***.

The hardware inventory report is downloaded as a .csv file with complete information as of the current date.

4. Provide this information to your NetApp support contact for contract renewal.

Viewing the Storage VM list and details

From the Health: All Storage VMs view, you can monitor your inventory of storage virtual machines (SVMs). You can use the Storage VM / Health details page to view detailed information about SVMs that are monitored.

What you’ll need

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view SVM details, such as the capacity, efficiency, and configuration of an SVM. You can also view information about the related devices and related alerts for that SVM.

Steps

1. In the left navigation pane, click **Storage > Storage VMs**.
2. Choose one of the following ways to view the SVM details:
 - To view information about the health of all SVMs in all clusters, in the View menu, select Health: All Storage VMs view.
 - To view the complete details, click the Storage VM name.

You can also view the complete details by clicking **View Details** in the minimal details dialog box.

3. View the objects related to the SVM by clicking **View Related** in the minimal details dialog box.

Related information

- [Storage VM: Health details page](#)
- [Performance: All Storage VMs view](#)
- [Security: Anti-ransomware view](#)
- [Viewing security status for clusters and Storage VMs](#)
- [Relationship: All Relationships view](#)

Viewing the aggregate list and details

From the Health: All Aggregates view, you can monitor your inventory of aggregates. The Capacity: All Aggregates view enables you to view information about the capacity and utilization of aggregates in all clusters.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

You can view details such as aggregate capacity and configuration, and disk information from the Aggregate / Health details page. You can use these details before you configure the threshold settings if required.

Steps

1. In the left navigation pane, click **Storage > Aggregates**.
2. Choose one of the following ways to view the aggregate details:
 - To view information about the health of all aggregates in all clusters, in the View menu, select Health: All Aggregates view.
 - To view information about the capacity and utilization of all aggregates in all clusters, in the View menu, select Capacity: All Aggregates view.
 - To view the complete details, click the aggregate name.

You can also view the complete details by clicking **View Details** in the minimal details dialog box.

3. View the objects related to the aggregate by clicking **View Related** from the minimal details dialog box.

Related information

- [Aggregate / Health details page](#)
- [Performance: All Aggregates view](#)
- [Customizing aggregate capacity reports](#)

Viewing FabricPool capacity information

You can view FabricPool capacity information for clusters, aggregates, and volumes on the Capacity and Performance inventory and details pages for these objects. These pages also display FabricPool mirror information when a mirror tier has been configured.

These pages display information such as the available capacity on the local performance tier and on the cloud tier, how much capacity is being used in both tiers, which aggregates are attached to a cloud tier, and which volumes are implementing the FabricPool features by moving certain information to the cloud tier.

When a cloud tier is mirrored to another cloud provider (the “mirror tier”) then both cloud tiers are displayed in the Aggregate / Health details page.

Steps

1. Perform one of the following:

To view capacity information for...	Do this...
Clusters	<ol style="list-style-type: none">a. On the Capacity: All Clusters view, click a cluster.b. On the Cluster / Health details page, click the Configuration tab. <p>The display shows the names of any cloud tiers to which this cluster is connected.</p>

To view capacity information for...	Do this...
Aggregates	<p>a. On the Capacity: All Aggregates view, click an aggregate where the Type field indicates “SSD (FabricPool)” or “HDD (FabricPool)”.</p> <p>b. On the Aggregate / Health details page, click the Capacity tab.</p> <p>The display shows the total capacity used in the cloud tier.</p> <p>c. Click the Disk Information tab.</p> <p>The display shows the name of the cloud tier and the capacity used.</p> <p>d. Click the Configuration tab.</p> <p>The display shows the name of the cloud tier and other detailed information about the cloud tier.</p>
Volumes	<p>a. On the Capacity: All Volumes view, click a volume where a policy name appears in the “Tiering Policy” field.</p> <p>b. On the Volume / Health details page, click the Configuration tab.</p> <p>The display shows the name of the FabricPool tiering policy assigned to the volume.</p>

- In the **Workload Analysis** page you can select “Cloud Tier View” in the **Capacity Trend** area to see the capacity being used in the local Performance Tier and in the Cloud Tier over the previous month.

For more information on FabricPool aggregates, see [Disks and aggregates overview](#).

Viewing storage pool details

You can view the details of the storage pool to monitor the storage pool health, total and available cache, and used and available allocations.

What you’ll need

You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

- In the left navigation pane, click **Storage > Aggregates**.
- Click an aggregate name.

The details of the selected aggregate are displayed.

3. Click the **Disk Information** tab.

Detailed disk information is displayed.



The Cache table is displayed only when the selected aggregate is using a storage pool.

4. In the Cache table, move the pointer over the name of the required storage pool.

The details of the storage pool are displayed.

Viewing the volume list and details

From the Health: All Volumes view, you can monitor your inventory of volumes. The Capacity: All Volumes view enables you to view information about the capacity and utilization of volumes in a cluster.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

You can also use the Volume / Health details page to view detailed information about volumes that are monitored, including the capacity, efficiency, configuration, and protection of the volumes. You can also view information about the related devices and related alerts for a specific volume.

Steps

1. In the left navigation pane, click **Storage > Volumes**.
2. Choose one of the following ways to view the volume details:
 - To view detailed information about the health of volumes in a cluster, in the View menu, select Health: All Volumes view.
 - To view detailed information about the capacity and utilization of volumes in a cluster, in the View menu, select Capacity: All Volumes view.
 - To view the complete details, click the volume name.

You can also view the complete details by clicking **View Details** in the minimal details dialog box.

3. **Optional:** View the objects related to the volume by clicking **View Related** from the minimal details dialog box.

Related information

- [Volume: Health details page](#)
- [Performance: All Volumes view](#)
- [Security: Anti-ransomware view](#)
- [Viewing volume protection relationships](#)
- [Creating a report to view available volume capacity charts](#)

Viewing details about NFS shares

You can view details about all NFS shares, such as its status, the path associated with

the volume (FlexGroup volumes or FlexVol volumes), access levels of clients to the NFS shares, and the export policy defined for the volumes that are exported. Use the Health: All NFS Shares view to see all NFS shares on all monitored clusters, and use the Storage VM / Health details page to view all NFS shares on a specific storage virtual machine (SVM).

What you'll need

- NFS license must be enabled on the cluster.
- Network interfaces serving the NFS shares must be configured.
- You must have the Operator, Application Administrator, or Storage Administrator role.

Step

1. In the left navigation pane, follow the steps below depending on whether you want to view all NFS shares or just the NFS shares for a particular SVM.

To...	Follow these steps...
View all NFS shares	Click Storage > NFS Shares
View NFS shares for single SVM	<ol style="list-style-type: none">a. Click Storage > Storage VMsb. Click the SVM for which you want to view the NFS shares details.c. In the Storage VM / Health details page, click the NFS Shares tab.

For more information, see [Provisioning file share volumes](#) and [Provisioning CIFS and NFS file shares by using APIs](#).

Viewing details about SMB/CIFS shares

You can view details about all SMB/CIFS shares, such as the share name, junction path, containing objects, security settings, and export policies defined for the share. Use the Health: All SMB Shares view to see all SMB shares on all monitored clusters, and use the Storage VM / Health details page to view all SMB shares on a specific storage virtual machine (SVM).

What you'll need

- CIFS license must be enabled on the cluster.
- Network interfaces serving the SMB/CIFS shares must be configured.
- You must have the Operator, Application Administrator, or Storage Administrator role.



Shares in folders are not displayed.

Step

1. In the left navigation pane, follow the steps below depending on whether you want to view all SMB/CIFS

shares or just the shares for a particular SVM.

To...	Follow these steps...
View all SMB/CIFS shares	Click Storage > SMB Shares
View SMB/CIFS shares for single SVM	<ol style="list-style-type: none">a. Click Storage > Storage VMsb. Click the SVM for which you want to view the SMB/CIFS share details.c. In the Storage VM / Health details page, click the SMB Shares tab.

For more information, see [Provisioning CIFS and NFS file shares by using APIs](#).

Viewing the list of Snapshot copies

You can view the list of Snapshot copies for a selected volume. You can use the list of Snapshot copies to calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted, and you can delete the Snapshot copies if required.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- The volume containing the Snapshot copies must be online.

Steps

1. In the left navigation pane, click **Storage > Volumes**.
2. In the **Health: All Volumes** view, select the volume that contains the Snapshot copies you want to view.
3. In the **Volume / Health** details page, click the **Capacity** tab.
4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

The number of Snapshot copies is a link that displays the list of Snapshot copies.

Related information

[Health/Volumes page](#)

Deleting Snapshot copies

You can delete a Snapshot copy to conserve space or to free disk space, or you can delete the Snapshot copy if it is no longer required.

What you'll need

You must have the Application Administrator or Storage Administrator role.

The volume must be online.

To delete a Snapshot copy that is busy or locked, you must have released the Snapshot copy from the application that was using it.

- You cannot delete the base Snapshot copy in a parent volume if a FlexClone volume is using that Snapshot copy.

The base Snapshot copy is the Snapshot copy that is used to create the FlexClone volume and displays the status `Busy` and Application Dependency as `Busy, Vclone` in the parent volume.

- You cannot delete a locked Snapshot copy that is used in a SnapMirror relationship.

The Snapshot copy is locked and is required for the next update.

Steps

1. In the left navigation pane, click **Storage > Volumes**.
2. In the **Health: All Volumes** view, select the volume that contains the Snapshot copies you want to view.

The list of Snapshot copies is displayed.
3. In the **Volume / Health** details page, click the **Capacity** tab.
4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

The number of Snapshot copies is a link that displays the list of Snapshot copies.
5. In the **Snapshot Copies** view, select the Snapshot copies you want to delete, and then click **Delete Selected**.

Calculating reclaimable space for Snapshot copies

You can calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- The volume must be online.
- The volume must be a FlexVol volume; this capability is not supported with FlexGroup volumes.

Steps

1. In the left navigation pane, click **Storage > Volumes**.
2. In the **Health: All Volumes** view, select the volume that contains the Snapshot copies you want to view.

The list of Snapshot copies is displayed.
3. In the **Volume / Health** details page, click the **Capacity** tab.
4. In **Details** pane of the **Capacity** tab, in the Other Details section, click the link next to **Snapshot Copies**.

The number of Snapshot copies is a link that displays the list of Snapshot copies.
5. In the **Snapshot Copies** view, select the Snapshot copies for which you want to calculate the reclaimable space.

6. Click **Calculate**.

The reclaimable space (in percentage, and KB, MB, GB, and so on) on the volume is displayed.

7. To recalculate the reclaimable space, select the required Snapshot copies and click **Recalculate**.

Description of cluster object windows and dialog boxes

You can view all your clusters and cluster objects from the respective storage object page. You can also view the details from the corresponding storage object details page. You can now launch System Manager user interface from the following STORAGE and PROTECTION sections of the INVENTORY.

- Cluster Inventory, Cluster Health, and Cluster Performance pages
- Aggregate Inventory, Aggregate Health, and Aggregate Performance pages
- Volume Inventory, Volume Health, and Volume Performance pages
- Node Inventory and Node Performance pages
- StorageVM Inventory, StorageVM Health, and StorageVM Performance pages
- Protection relationship pages

Common Unified Manager health workflows and tasks

Some common administrative workflows and tasks associated with Unified Manager include selecting the storage clusters that are to be monitored; diagnosing conditions that adversely affect data availability, capacity, and protection; restoring lost data; configuring and managing volumes; and bundling and sending diagnostic data to technical support (when necessary).

Unified Manager enables storage administrators to view a dashboard, assess the overall capacity, availability, and protection health of the managed storage clusters, and then quickly identify, locate, diagnose, and assign for resolution any specific issues that might arise.

The most important issues related to a cluster, storage virtual machine (SVM), volume, or FlexGroup volume that affect the storage capacity or data availability of your managed storage objects are displayed in the system health graphs and events on the Dashboard page. When critical issues are identified, this page provides links to support appropriate troubleshooting workflows.

Unified Manager can also be included in workflows that include related manageability tools—such as OnCommand Workflow Automation (WFA)—to support the direct configuration of storage resources.

Common workflows related to the following administrative tasks are described in this document:

- Diagnosing and managing availability issues

If hardware failure or storage resource configuration issues cause the display of data availability events in the Dashboard page, storage administrators can follow the embedded links to view connectivity information about the affected storage resource, view troubleshooting advice, and assign issue resolution to other administrators.

- Configuring and monitoring performance incidents

The Administrator can monitor and manage the performance of the storage system resources that are being monitored. See the [Introduction to Active IQ Unified Manager performance monitoring](#) for more information.

- Diagnosing and managing volume capacity issues

If volume storage capacity issues are displayed in the Dashboard page, storage administrators can follow the embedded links to view the current and historical trends related to the storage capacity of the affected volume, view troubleshooting advice, and assign issue resolution to other administrators.

- Configuring, monitoring, and diagnosing protection relationship issues

After creating and configuring protection relationships, storage administrators can view the potential issues related to protection relationships, the current state of the protection relationships, the current and historical protection job success information about the affected relationships, and troubleshooting advice. See the [Creating, monitoring, and troubleshooting protection relationships](#) for more information.

- Creating backup files and restoring data from backup files.

- Associating storage objects with annotations

By associating storage objects with annotations, storage administrators can filter and view the events that are related to the storage objects, which enables storage administrators to prioritize and resolve the issues that are associated with the events.

- Using REST APIs to help manage your clusters by viewing the health, capacity, and performance information captured by Unified Manager. See [Getting started with Active IQ Unified Manager REST APIs](#) for more information.

- Sending a support bundle to technical support

Storage administrators can retrieve and send a support bundle to technical support by using the maintenance console. Support bundles must be sent to technical support when the issue requires more detailed diagnosis and troubleshooting than what an AutoSupport message provides.

Monitoring and troubleshooting data availability

Unified Manager monitors the reliability with which authorized users can access your stored data, alerts you to conditions that block or impede that access, and enables you to diagnose those conditions and assign and track their resolution.

The availability workflow topics in this section describe examples of how a storage administrator can use the Unified Manager web UI to discover, diagnose, and assign for resolution hardware and software conditions that adversely affect data availability.

Scanning for and resolving storage failover interconnect link down conditions

This workflow provides an example of how you might scan for, evaluate, and resolve downed storage failover interconnect link conditions. In this scenario, you are an administrator using Unified Manager to scan for storage failover risks before starting an ONTAP version upgrade on your nodes.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

If storage failover interconnections between HA pair nodes fail during a nondisruptive upgrade attempt, the upgrade fails. Therefore, common practice is for the administrator to monitor and confirm storage failover reliability on the cluster nodes targeted for upgrade before the start of an upgrade.

Steps

1. In the left navigation pane, click **Event Management**.
2. In the **Event Management** inventory page, select **Active Availability events**.
3. At the top of the **Event Management** inventory page **Name** column, click  and enter `*failover` in the text box to limit the event to display to storage failover-related events.

All past events related to storage failover conditions are displayed.

In this scenario, the Unified Manager displays the event, "Storage Failover Interconnect One or More Links Down" in its Availability Incidents section.

4. If one or more events related to storage failover are displayed on the **Event Management** inventory page, perform the following steps:
 - a. Click the event title link to display event details for that event.

In this example, you click the event title "Storage Failover Interconnect One or More Links Down".

The Event details page for that event is displayed.

- b. On the Event details page, you can perform one or more of the following tasks:
 - Review the error message in the Cause field and evaluate the issue.
 - Assign the event to an administrator.
 - Acknowledge the event.

Related information

[Event details page](#)

[Unified Manager user roles and capabilities](#)

Performing corrective action for storage failover interconnect links down

When you display the Event details page of a storage failover-related event, you can review the summary information of the page to determine the urgency of the event, possible cause of the issue, and possible resolution to the issue.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

In this example scenario, the event summary provided on the Event details page contains the following information about the storage failover interconnect link down condition:

Event: Storage Failover Interconnect One or More Links Down

Summary

Severity: Warning

State: New

Impact Level: Risk

Impact Area: Availability

Source: aardvark

Source Type: Node

Acknowledged By:

Resolved By:

Assigned To:

Cause: At least one storage failover interconnected link between the nodes aardvark and bonobo is down. RDMA interconnect is up (Link0 up, Link1 down)

The example event information indicates that a storage failover interconnect link, Link1, between HA pair nodes aardvark and bonobo is down, but that Link0 between Apple and Boy is active. Because one link is active, the remote dynamic memory access (RDMA) is still functioning and a storage failover job can still succeed.

However, to ensure against both links failing and storage failover protection being totally disabled, you decide to further diagnose the reason for Link1 going down.

Steps

1. From the **Event** details page, you can click the link to the event specified in the Source field to obtain further details of other events that might be related to the storage failover interconnection link down condition.

In this example, the source of the event is the node named aardvark. Clicking that node name displays the HA Details for the affected HA pair, aardvark and bonobo, on the Nodes tab of the Cluster / Health details page, and displays other events that recently occurred on the affected HA pair.

2. Review the **HA Details** for more information relating to the event.

In this example, the relevant information is in the Events table. The table shows the “Storage Failover Connection One or More Link Down” event, the time the event was generated, and, again, the node from which this event originated.

Using the node location information in the HA Details, request or personally complete a physical inspection and repair of the storage failover issue on the affected HA pair nodes.

Related information

[Event details page](#)

[Unified Manager user roles and capabilities](#)

Resolving volume offline issues

This workflow provides an example of how you might evaluate and resolve a volume offline event that Unified Manager might display in the Event Management inventory page. In this scenario, you are an administrator using Unified Manager to troubleshoot one or more volume offline events.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

Volumes might be reported offline for several reasons:

- The SVM administrator has deliberately taken the volume offline.
- The volume's hosting cluster node is down and storage failover to its HA pair partner has failed also.
- The volume's hosting storage virtual machine (SVM) is stopped because the node hosting the root volume of that SVM is down.
- The volume's hosting aggregate is down due to simultaneous failure of two RAID disks.

You can use the Event Management inventory page and the Cluster/Health, Storage VM/Health, and Volume/Health details pages to confirm or eliminate one or more of these possibilities.

Steps

1. In the left navigation pane, click **Event Management**.
2. In the **Event Management** inventory page, select **Active Availability events**.
3. Click the hypertext link displayed for the Volume Offline event.

The Event details page for the availability incident is displayed.

4. On that page, check the notes for any indication that the SVM administrator has taken the volume in question offline.
5. On the **Event** details page, you can review the information for one or more of the following tasks:
 - Review the information displayed in the Cause field for possible diagnostic guidance.

In this example, the information in the Cause field informs you only that the volume is offline.

- Check the Notes and Updates area for any indication that the SVM administrator has deliberately taken the volume in question offline.
- Click the source of the event, in this case the volume that is reported offline, to get more information about that volume.
- Assign the event to an administrator.
- Acknowledge the event or, if appropriate, mark it as resolved.

Performing diagnostic actions for volume offline conditions

After navigating to the Volume / Health details page of a volume reported to be offline, you can search for additional information helpful to diagnosing the volume offline condition.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

If the volume that is reported offline was not taken offline deliberately, that volume might be offline for several reasons.

Starting at the offline volume's Volume / Health details page, you can navigate to other pages and panes to confirm or eliminate possible causes:

- Click **Volume / Health** details page links to determine if the volume is offline because its host node is down and storage failover to its HA pair partner has failed also.

See [Determining if a volume offline condition is caused by a down node](#).

- Click **Volume / Health** details page links to determine if the volume is offline and its host storage virtual machine (SVM) is stopped because the node hosting the root volume of that SVM is down.

See [Determining if a volume is offline and SVM is stopped because a node is down](#).

- Click **Volume / Health** details page links to determine if the volume is offline because of broken disks in its host aggregate.

See [Determining if a volume is offline because of broken disks in an aggregate](#).

Related information

[Unified Manager user roles and capabilities](#)

Determining if a volume is offline because its host node is down

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because its host node is down and that storage failover to its HA pair partner is unsuccessful.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

To determine if the volume offline condition is caused by failure of the hosting node and subsequent unsuccessful storage failover, perform the following actions:

Steps

1. Locate and click the hypertext link displayed under SVM in the **Related Devices** pane of the offline volume's **Volume / Health** details page.

The Storage VM / Health details page displays information about the offline volume's hosting storage virtual machine (SVM).

2. In the **Related Devices** pane of the **Storage VM / Health** details page, locate and click hypertext link displayed under Volumes.

The Health: All Volumes view displays a table of information about all the volumes hosted by the SVM.

3. On the **Health: All Volumes** view State column header, click the filter symbol , and then select the option **Offline**.

Only the SVM volumes that are in offline state are listed.

4. On the Health: All Volumes view, click the grid symbol , and then select the option **Cluster Nodes**.

You might need to scroll in the grid selection box to locate the **Cluster Nodes** option.

The Cluster Nodes column is added to the volumes inventory and displays the name of the node that hosts each offline volume.

5. On the **Health: All Volumes** view, locate the listing for the offline volume and, in its Cluster Node column, click the name of its hosting node.

The Nodes tab on the Cluster / Health details page displays the state of the HA pair of nodes to which the hosting node belongs. The state of the hosting node and the success of any cluster failover operation is indicated in the display.

After you confirm that the volume offline condition exists because its host node is down and storage failover to the HA pair partner has failed, contact the appropriate administrator or operator to manually restart the down node and fix the storage failover problem.

Determining if a volume is offline and its SVM is stopped because a node is down

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because its host storage virtual machine (SVM) is stopped due to the node hosting the root volume of that SVM being down.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

To determine if the volume offline condition is caused its host SVM being stopped because the node hosting the root volume of that SVM is down, perform the following actions:

Steps

1. Locate and click the hypertext link displayed under the SVM in the **Related Devices** pane of the offline volume's **Volume / Health** details page.

The Storage VM / Health details page displays the "running" or the "stopped" status of the hosting SVM. If the SVM status is running, then the volume offline condition is not caused by the node hosting the root volume of that SVM being down.

2. If the SVM status is stopped, then click **View SVMs** to further identify the cause of the hosting SVM being stopped.
3. On the **Health: All Storage VMs** view SVM column header, click the filter symbol  and then type the name of the stopped SVM.

The information for that SVM is shown in a table.

4. On the **Health: All Storage VMs** view, click  and then select the option **Root Volume**.

The Root Volume column is added to the SVM inventory and displays the name of the root volume of the

stopped SVM.

5. In the Root Volume column, click the name of the root volume to display the **Storage VM / Health** details page for that volume.

If the status of the SVM root volume is (Online), then the original volume offline condition is not caused because the node hosting the root volume of that SVM is down.

6. If the status of the SVM root volume is (Offline), then locate and click the hypertext link displayed under Aggregate in the Related Devices pane of the SVM root volume's Volume / Health details page.
7. Locate and click the hypertext link displayed under Node in the **Related Devices** pane of the Aggregate's **Aggregate / Health** details page.

The Nodes tab on the Cluster / Health details page displays the state of the HA pair of nodes to which the SVM root volume's hosting node belongs. The state of the node is indicated in the display.

After you confirm that the volume offline condition is caused by that volume's host SVM offline condition, which itself is caused by the node that hosts the root volume of that SVM being down, contact the appropriate administrator or operator to manually restart the down node.

Determining if a volume is offline because of broken disks in an aggregate

You can use the Unified Manager web UI to confirm or eliminate the possibility that a volume is offline because RAID disk problems have taken its host aggregate offline.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

To determine if the volume offline condition is caused by RAID disk problems that are taking the hosting aggregate offline, perform the following actions:

Steps

1. Locate and click the hypertext link displayed under Aggregate in the **Related Devices** pane of the **Volume / Health** details page.

The Aggregate / Health details page displays the online or offline status of the hosting aggregate. If the aggregate status is online, then RAID disk problems are not the cause of the volume being offline.

2. If the aggregate status is offline, then click **Disk Information** and look for broken disk events in the **Events** list on the **Disk Information** tab.
3. To further identify the broken disks, click the hypertext link displayed under Node in the **Related Devices** pane.

The Cluster / Health details page is displayed.

4. Click **Disks**, and then select **Broken** in the **Filters** pane to list all disks in the broken state.

If the disks in the broken state caused the offline state of the host aggregate, the name of the aggregate is displayed in the Impacted Aggregate column.

After confirming that the volume offline condition is caused by broken RAID disks and the consequent offline host aggregate, contact the appropriate administrator or operator to manually replace the broken disks and put

the aggregate back online.

Resolving capacity issues

This workflow provides an example of how you can resolve a capacity issue. In this scenario, you are an administrator or operator and you access the Unified ManagerDashboard page to see if any of the monitored storage objects have capacity issues. You want to determine the possible cause of and resolution to the problem.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

On the Dashboard page, you look for a "Volume Space Full" error event in the Capacity panel under the events drop-down list.

Steps

1. In the **Capacity** panel of the **Dashboard** page, click the name of the Volume Space Full error event.

The Event details page for the error is displayed.

2. From the **Event** details page, you can perform one or more of the following tasks:
 - Review the error message in the Cause field and click the suggestions under Suggested Remedial Actions to review descriptions of possible remediations.
 - Click the object name, in this case a volume, in the Source field to get details about the object.
 - Look for notes that might have been added about this event.
 - Add a note to the event.
 - Assign the event to another user.
 - Acknowledge the event.
 - Mark the event as resolved.

Related information

[Event details page](#)

Performing suggested remedial actions for a full volume

After receiving a "Volume Space Full" error event, you review the suggested remedial actions on the Event details page and decide to perform one of the suggested actions.

What you'll need

You must have the Application Administrator or Storage Administrator role.

A user with any role can perform all of the tasks in this workflow that use Unified Manager.

In this example, you have seen a Volume Space Full error event on the Unified ManagerEvent Management inventory page and have clicked the name of the event.

Possible remedial actions you might perform for a full volume include the following:

- Enabling autogrow, deduplication, or compression on the volume
- Resizing or moving the volume
- Deleting or moving data from the volume

Although all of these actions must be performed from either ONTAP System Manager or the ONTAP CLI, you can use Unified Manager to find information you might need to determine which actions to take.

Steps

1. From the **Event** details page, you click the volume name in the Source field to view details about the affected volume.
2. On the **Volume / Health** details page, you click **Configuration** and see that deduplication and compression are already enabled on the volume.

You decide to resize the volume.

3. In the **Related Devices** pane, you click the name of the hosting aggregate to see if the aggregate can accommodate a larger volume.
4. On the **Aggregate / Health** details page, you see that the aggregate hosting the full volume does have enough uncommitted capacity, so you use ONTAP System Manager to resize the volume, giving it more capacity.

Related information

[Event details page](#)

Managing health thresholds

You can configure global health threshold values for all the aggregates, volumes, and qtrees to track any health threshold breaches.

What storage capacity health thresholds are

A storage capacity health threshold is the point at which the Unified Manager server generates events to report any capacity problem with storage objects. You can configure alerts to send notification whenever such events occurs.

The storage capacity health thresholds for all aggregates, volumes, and qtrees are set to default values. You can change the settings as required for an object or a group of objects.

Configuring global health threshold settings

You can configure global health threshold conditions for capacity, growth, Snapshot reserve, quotas, and inodes to monitor your aggregate, volume, and qtree size effectively. You can also edit the settings for generating events for exceeding lag thresholds.

Global health threshold settings apply to all objects with which they are associated, such as aggregates, volumes, and so forth. When thresholds are crossed, an event is generated and, if alerts are configured, an alert notification is sent. Threshold defaults are set to recommended values, but you can modify them to generate events at intervals to meet your specific needs. When thresholds are changed, events are generated or obsoleted in the next monitoring cycle.

Global health threshold settings are accessible from the Event Thresholds section of the left-navigation menu. You can also modify threshold settings for individual objects, from the inventory page or the details page for that object.

- For information, see [Configuring global aggregate health threshold values](#).

You can configure the health threshold settings for capacity, growth, and Snapshot copies for all aggregates to track any threshold breach.

- For information, see [Configuring global volume health threshold values](#).

You can edit the health threshold settings for capacity, Snapshot copies, qtree quotas, volume growth, overwrite reserve space, and inodes for all volumes to track any threshold breach.

- For information, see [Configuring global qtree health threshold values](#).

You can edit the health threshold settings for capacity for all qtrees to track any threshold breach.

- For information, see [Editing lag health threshold settings for unmanaged protection relationships](#).

You can increase or decrease the warning or error lag time percentage so that events are generated at intervals that are more appropriate to your needs.

Configuring global aggregate health threshold values

You can configure global health threshold values for all aggregates to track any threshold breach. Appropriate events are generated for threshold breaches and you can take preventive measures based on these events. You can configure the global values based on the best practice settings for thresholds that apply to all monitored aggregates.

What you'll need

You must have the Application Administrator or Storage Administrator role.

When you configure the options globally, the default values of the objects are modified. However, if the default values have been changed at the object level, the global values are not modified.

The threshold options have default values for better monitoring, however, you can change the values to suit the requirements of your environment.

When Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.



Health threshold values are not applicable to the root aggregate of the node.

Steps

1. In the left navigation pane, click **Event Thresholds > Aggregate**.
2. Configure the appropriate threshold values for capacity, growth, and Snapshot copies.
3. Click **Save**.

Related information

[Adding users](#)

Configuring global volume health threshold values

You can configure the global health threshold values for all volumes to track any threshold breach. Appropriate events are generated for health threshold breaches, and you can take preventive measures based on these events. You can configure the global values based on the best practice settings for thresholds that apply to all monitored volumes.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Most of the threshold options have default values for better monitoring. However, you can change the values to suit the requirements of your environment.

Note that when Autogrow is enabled on a volume that capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.



The default value of 1000 Snapshot copies is applicable only to FlexVol volumes when the ONTAP version is 9.4 or greater, and to FlexGroup volumes when the ONTAP version is 9.8 and greater. For clusters installed with older versions of ONTAP software, the maximum number is 250 Snapshot copies per volume. For these older versions, Unified Manager interprets this number 1000 (and any number between 1000 and 250) as 250; meaning you will continue to receive events when the number of Snapshot copies reaches 250. If you wish to set this threshold to less than 250 for these older versions, you must set the threshold to 250 or lower here, in the Health: All Volumes view, or in the Volume / Health details page.

Steps

1. In the left navigation pane, click **Event Thresholds > Volume**.
2. Configure the appropriate threshold values for capacity, Snapshot copies, qtree quotas, volume growth, and inodes.
3. Click **Save**.

Related information

[Adding users](#)

Configuring global qtree health threshold values

You can configure the global health threshold values for all qtrees to track any threshold breach. Appropriate events are generated for health threshold breaches, and you can take preventive measures based on these events. You can configure the global values based on the best practice settings for thresholds that apply to all monitored qtrees.

What you'll need

You must have the Application Administrator or Storage Administrator role.

The threshold options have default values for better monitoring, however, you can change the values to suit the requirements of your environment.

Events are generated for a qtree only when a Qtree quota or a Default quota has been set on the qtree. Events are not generated if the space defined in a User quota or Group quota has exceeded the threshold.

Steps

1. In the left navigation pane, click **Event Thresholds > Qtree**.
2. Configure the appropriate capacity threshold values.
3. Click **Save**.

Configuring lag threshold settings for unmanaged protection relationships

You can edit the global default lag warning and error health threshold settings for unmanaged protection relationships so that events are generated at intervals that are appropriate to your needs.

What you'll need

You must have the Application Administrator or Storage Administrator role.

The lag time must be no more than the defined transfer schedule interval. For example, if the transfer schedule is hourly, then the lag time must not be more than one hour. The lag threshold specifies a percentage that the lag time must not exceed. Using the example of one hour, if the lag threshold is defined as 150%, then you will receive an event when the lag time is more than 1.5 hours.

The settings described in this task are applied globally to all unmanaged protection relationships. The settings cannot be specified and applied exclusively to one unmanaged protection relationship.

Steps

1. In the left navigation pane, click **Event Thresholds > Relationship**.
2. Increase or decrease the global default warning or error lag time percentage as required.
3. To disable a warning or error event from being triggered from any lag threshold amount, uncheck the box next to **Enabled**.
4. Click **Save**.

Related information

[Adding users](#)

Editing individual aggregate health threshold settings

You can edit the health threshold settings for aggregate capacity, growth, and Snapshot copies of one or more aggregates. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

When Autogrow is enabled on volumes that reside on the aggregate, the aggregate capacity thresholds are

considered breached based on the maximum volume size set by autogrow, not based on the original volume size.

Steps

1. In the left navigation pane, click **Storage > Aggregates**.
2. In the **Health: All Aggregates** view, select one or more aggregates and then click **Edit Thresholds**.
3. In the **Edit Aggregate Thresholds** dialog box, edit the threshold settings of one of the following: capacity, growth, or Snapshot copies by selecting the appropriate check box and then modifying the settings.
4. Click **Save**.

Related information

[Adding users](#)

Editing individual volume health threshold settings

You can edit the health threshold settings for volume capacity, growth, quota, and space reserve of one or more volumes. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

Note that when Autogrow is enabled on a volume that capacity thresholds are considered breached based on the maximum volume size set by autogrow, not based on the original volume size.



The default value of 1000 Snapshot copies is applicable only to FlexVol volumes when the ONTAP version is 9.4 or greater, and to FlexGroup volumes when the ONTAP version is 9.8 and greater. For clusters installed with older versions of ONTAP software, the maximum number is 250 Snapshot copies per volume. For these older versions, Unified Manager interprets this number 1000 (and any number between 1000 and 250) as 250; meaning you will continue to receive events when the number of Snapshot copies reaches 250. If you wish to set this threshold to less than 250 for these older versions, you must set the threshold to 250 or lower here, in the Health: All Volumes view, or in the Volume / Health details page.

Steps

1. In the left navigation pane, click **Storage > Volumes**.
2. In the **Health: All Volumes** view, select one or more volumes and then click **Edit Thresholds**.
3. In the **Edit Volume Thresholds** dialog box, edit the threshold settings of one of the following: capacity, Snapshot copies, qtree quota, growth, or inodes by selecting the appropriate check box and then modifying the settings.
4. Click **Save**.

Related information

[Adding users](#)

Editing individual qtree health threshold settings

You can edit the health threshold settings for qtree capacity for one or more qtrees. When a threshold is crossed, alerts are generated and you receive notifications. These notifications help you to take preventive measures based on the event generated.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Based on changes to the threshold values, events are generated or obsoleted in the next monitoring cycle.

Steps

1. In the left navigation pane, click **Storage > Qtrees**.
2. In the **Capacity: All Qtrees** view, select one or more qtrees and then click **Edit Thresholds**.
3. In the **Edit Qtree Thresholds** dialog box, change the capacity thresholds for the selected qtree or qtrees and click **Save**.



You can also set individual qtree thresholds from the Qtrees tab on the Storage VM / Health details page.

Managing cluster security objectives

Unified Manager provides a dashboard that identifies how secure your ONTAP clusters, storage virtual machines (SVMs), and volumes are based on recommendations defined in the *NetApp Security Hardening Guide for ONTAP 9*.

The goal of the security dashboard is to show any areas where your ONTAP clusters do not align with the NetApp recommended guidelines so that you can fix these potential issues. In most cases you will fix the issues using ONTAP System Manager or the ONTAP CLI. Your organization may not follow all of the recommendations, so in some cases you will not need to make any changes.

See the [NetApp Security Hardening Guide for ONTAP 9 \(TR-4569\)](#) for detailed recommendations and resolutions.

In addition to reporting security status, Unified Manager also generates security events for any cluster or SVM that has security violations. You can track these issues in the Event Management inventory page and you can configure alerts for these events so that your storage administrator is notified when new security events occur.

For more information, see [What security criteria are being evaluated](#).

What security criteria are being evaluated

In general, security criteria for your ONTAP clusters, storage virtual machines (SVMs), and volumes are being evaluated against the recommendations defined in the *NetApp Security Hardening Guide for ONTAP 9*.

Some of the security checks include:

- whether a cluster is using a secure authentication method, such as SAML

- whether peered clusters have their communication encrypted
- whether a storage VM has its audit log enabled
- whether your volumes have software or hardware encryption enabled

See the topics on compliance categories and the [NetApp Security Hardening Guide for ONTAP 9](#) for detailed information.



Upgrade events that are reported from the Active IQ platform are also considered security events. These events identify issues where the resolution requires you to upgrade ONTAP software, node firmware, or operating system software (for security advisories). These events are not displayed in the Security panel, but they are available from the Event Management inventory page.

For more information, see [Managing cluster security objectives](#).

Cluster compliance categories

This table describes the cluster security compliance parameters that Unified Manager evaluates, the NetApp recommendation, and whether the parameter affects the overall determination of the cluster being compliant or not compliant.

Having non-compliant SVMs on a cluster will affect the compliance value for the cluster. So in some cases you may need to fix a security issues with an SVM before your cluster security is seen as compliant.

Note that not every parameter listed below appears for all installations. For example, if you have no peered clusters, or if you have disabled AutoSupport on a cluster, then you will not see the Cluster Peering or AutoSupport HTTPS Transport items in the UI page.

Parameter	Description	Recommendation	Affects Cluster Compliance
Global FIPS	Indicates if Global FIPS (Federal Information Processing Standard) 140-2 compliance mode is enabled or disabled. When FIPS is enabled, TLSv1 and SSLv3 are disabled, and only TLSv1.1 and TLSv1.2 are allowed.	Enabled	Yes
Telnet	Indicates if Telnet access to the system is enabled or disabled. NetApp recommends Secure Shell (SSH) for secure remote access.	Disabled	Yes

Parameter	Description	Recommendation	Affects Cluster Compliance
Insecure SSH Settings	Indicates if SSH uses insecure ciphers, for example ciphers beginning with *cbc.	No	Yes
Login Banner	Indicates if the Login banner is enabled or disabled for users accessing the system.	Enabled	Yes
Cluster Peering	Indicates if communication between peered clusters is encrypted or unencrypted. Encryption must be configured on both the source and destination clusters for this parameter to be considered compliant.	Encrypted	Yes
Network Time Protocol	Indicates if the cluster has one or more configured NTP servers. For redundancy and best service NetApp recommends that you associate at least three NTP servers with the cluster.	Configured	Yes
OCSP	Indicates if there are applications in ONTAP that are not configured with OCSP (Online Certificate Status Protocol) and therefore communications are not encrypted. The non-compliant applications are listed.	Enabled	No
Remote Audit Logging	Indicates if log forwarding (Syslog) is encrypted or not encrypted.	Encrypted	Yes

Parameter	Description	Recommendation	Affects Cluster Compliance
AutoSupport HTTPS Transport	Indicates if HTTPS is used as the default transport protocol for sending AutoSupport messages to NetApp support.	Enabled	Yes
Default Admin User	Indicates if the Default Admin User (built-in) is enabled or disabled. NetApp recommends locking (disabling) any unneeded built-in accounts.	Disabled	Yes
SAML Users	Indicates if SAML is configured. SAML enables you to configure multi-factor authentication (MFA) as a login method for single sign-on.	No	No
Active Directory Users	Indicates if Active Directory is configured. Active Directory and LDAP are the preferred authentication mechanisms for users accessing clusters.	No	No
LDAP Users	Indicates if LDAP is configured. Active Directory and LDAP are the preferred authentication mechanisms for users managing clusters over local users.	No	No
Certificate Users	Indicates if a certificate user is configured to log into the cluster.	No	No
Local Users	Indicates if local users are configured to log into the cluster.	No	No

Parameter	Description	Recommendation	Affects Cluster Compliance
Remote Shell	Indicates if RSH is enabled. For security reasons, RSH should be disabled. The Secure Shell (SSH) for secure remote access is preferred.	Disabled	Yes
MD5 in Use	Indicates if ONTAP user accounts use less-secure MD5 Hash function. The MD5 Hashed user accounts migration to the more secure cryptographic hash function like SHA-512 is preferred.	No	Yes
Certificate Issuer Type	Indicates the type of digital certificate used.	CA-Signed	No

Storage VM compliance categories

This table describes the storage virtual machine (SVM) security compliance criteria that Unified Manager evaluates, the NetApp recommendation, and whether the parameter affects the overall determination of the SVM being complaint or not complaint.

Parameter	Description	Recommendation	Affects SVM Compliance
Audit Log	Indicates if Audit logging is enabled or disabled.	Enabled	Yes
Insecure SSH Settings	Indicates if SSH uses insecure ciphers, for example ciphers beginning with <code>cbc*</code> .	No	Yes
Login Banner	Indicates if the Login banner is enabled or disabled for users accessing SVMs on the system.	Enabled	Yes
LDAP Encryption	Indicates if LDAP Encryption is enabled or disabled.	Enabled	No

Parameter	Description	Recommendation	Affects SVM Compliance
NTLM Authentication	Indicates if NTLM Authentication is enabled or disabled.	Enabled	No
LDAP Payload Signing	Indicates if LDAP Payload Signing is enabled or disabled.	Enabled	No
CHAP Settings	Indicates if CHAP is enabled or disabled.	Enabled	No
Kerberos V5	Indicates if Kerberos V5 authentication is enabled or disabled.	Enabled	No
NIS Authentication	Indicates if the use of NIS authentication is configured.	Disabled	No
FPolicy Status Active	Indicates if FPolicy is created or not.	Yes	No
SMB Encryption Enabled	Indicates if SMB -Signing & Sealing is not enabled.	Yes	No
SMB Signing Enabled	Indicates if SMB -Signing is not enabled.	Yes	No

Volume compliance categories

This table describes the volume encryption parameters that Unified Manager evaluates to determine whether the data on your volumes is adequately protected from being accessed by unauthorized users.

Note that the volume encryption parameters do not affect whether the cluster or storage VM is considered compliant.

Parameter	Description
Software Encrypted	Displays the number of volumes that are protected using NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) software encryption solutions.
Hardware Encrypted	Displays the number of volumes that are protected using NetApp Storage Encryption (NSE) hardware encryption.

Parameter	Description
Software and Hardware Encrypted	Displays the number of volumes that are protected by both software and hardware encryption.
Not Encrypted	Displays the number of volumes that are not encrypted.

What does not compliant mean

Clusters and storage virtual machines (SVMs) are considered not compliant when any of the security criteria that is being evaluated against the recommendations defined in the *NetApp Security Hardening Guide for ONTAP 9* are not met. Additionally, a cluster is considered not compliant when any SVM is flagged as being not compliant.

The status icons in the security cards have the following meanings in relation to their compliance:

-  - The parameter is configured as recommended.
-  - The parameter is not configured as recommended.
-  - Either the functionality is not enabled on the cluster, or the parameter is not configured as recommended, but this parameter does not contribute to the compliance of the object.

Note that volume encryption status does not contribute to whether the cluster or SVM are considered compliant.

Viewing security status for clusters and Storage VMs

Active IQ Unified Manager enables you to view the security status of the storage objects in your environment from different points in the interface. You can collect and analyze information and reports based on defined parameters, and detect suspicious behavior or unauthorized system changes on the monitored clusters and storage VMs.

For the security recommendations, see the [NetApp Security Hardening Guide for ONTAP 9](#)

View object level security status on Security page

As a system administrator, you can use the **Security** page to get visibility into the security strength of your ONTAP clusters and storage VMs at the data center and site levels. The supported objects are cluster, storage VMs, and volumes. Follow these steps:

Steps

1. In the left navigation pane, click **Dashboard**.
2. Depending on whether you want to view security status for all monitored clusters or for a single cluster, select **All Clusters** or select a single cluster from the drop-down menu.
3. Click the right-arrow in the **Security** panel. The Security page is displayed.

Clicking the bar charts, counts, and [View Reports](#) links takes you to the Volumes, Clusters, or Storage VMs page for you to view the corresponding details or generate reports, as required.

The Security page displays the following panels:

- **Cluster Compliance:** the security status (number of clusters that are compliant or not compliant) of all the clusters in a data center
- **Storage VM Compliance:** the security status (number of storage VMs that are compliant or not compliant) for all the storage VMs in your data center
- **Volume Encryption:** the volume encryption status (number of volumes that are encrypted or not encrypted) of all the volumes in your environment
- **Volume Anti-ransomware Status:** the security status (number of volumes with anti-ransomware enabled or disabled) of all the volumes in your environment
- **Cluster Authentication and Certificates:** the number of clusters using each type of authentication method, such as SAML, Active Directory, or through certificates and local authentication. The panel also displays the number of clusters whose certificates have either expired or are about to expire in 60 days.

View security details of all clusters on the Clusters page

The **Clusters / Security** details page enables you to view the security compliance status at a cluster level.

Steps

1. In the left navigation pane, click **Storage > Clusters**.
2. Select **View > Security > All Clusters**.

Default security parameters, such as Global FIPS, Telnet, insecure SSH settings, login banner, network time protocol, AutoSupport HTTPS Transport, and the status of cluster certificate expiration are displayed.

You can click the  more options button and choose to view the security details on the **Security** page of Unified Manager or on System Manager. You should have valid credentials for viewing the details on System Manager.



If a cluster has an expired certificate, you can click `expired` under **Cluster Certificate Validity**, and renew it from System Manager (9.10.1 and later). You cannot click `expired` if the System Manager instance is of a release earlier than 9.10.1.

View security details of all clusters from the storage VMs page

The **Storage VMs / Security** details page enables you to view the security compliance status at a storage VM level.

Steps

1. In the left navigation pane, click **Storage > Storage VMs**.
2. Select **View > Security > All Storage VMs**. A list of clusters with the security parameters is displayed.

You can have a default view of the storage VMs' security compliance by checking the security parameters, such as storage VMs, cluster, login banner, audit log, and insecure SSH settings.

You can click the  more options button and choose to view the security details on the **Security** page of Unified Manager or on System Manager. You should have valid credentials for viewing the details on System Manager.

For anti-ransomware security details for volumes and storage VMs, see [Viewing the anti-ransomware status of all volumes and Storage VMs](#).

Viewing security events that may require software or firmware updates

There are certain security events that have an impact area of “Upgrade”. These events are reported from the Active IQ platform, and they identify issues where the resolution requires you to upgrade ONTAP software, node firmware, or operating system software (for security advisories).

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

You may want to perform immediate corrective action for some of these issues, whereas other issues may be able to wait until your next scheduled maintenance. You can view all of these events and assign them to users who can resolve the issues. Additionally, if there are certain security upgrade events that you do not want to be notified about, this list can help you identify those events so that you can disable them.

Steps

1. In the left navigation pane, click **Event Management**.

By default, all Active (New and Acknowledged) events are displayed on the Event Management inventory page.

2. From the View menu, select **Upgrade events**.

The page displays all active upgrade security events.

Viewing how user authentication is being managed on all clusters

The Security page displays the types of authentication being used to authenticate users on each cluster, and the number of users who are accessing the cluster using each type. This enables you to verify that user authentication is being performed securely as defined by your organization.

Steps

1. In the left navigation pane, click **Dashboard**.
2. At the top of the dashboard, select **All Clusters** from the drop-down menu.
3. Click the right-arrow in the **Security** panel and the **Security** page is displayed.
4. View the **Cluster Authentication** card to see the number of users who are accessing the system using each authentication type.
5. View the **Cluster Security** card to view the authentication mechanisms being used to authenticate users on each cluster.

If there are some users accessing the system using an insecure method, or using a method that is not recommended by NetApp, you can disable the method.

Viewing the encryption status of all volumes

You can view a list of all the volumes and their current encryption status so you can determine whether the data on your volumes is adequately protected from being accessed by unauthorized users.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

The types of encryption that can be applied to a volume are:

- Software - Volumes that are protected using NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) software encryption solutions.
- Hardware - Volumes that are protected using NetApp Storage Encryption (NSE) hardware encryption.
- Software and Hardware - Volumes that are protected by both software and hardware encryption.
- None - Volumes that are not encrypted.

Steps

1. In the left navigation pane, click **Storage > Volumes**.
2. In the View menu, select **Health > Volumes Encryption**
3. In the **Health: Volumes Encryption** view, sort on the **Encryption Type** field, or use the Filter to display volumes that have a specific encryption type, or that are not encrypted (Encryption Type of "None").

Viewing the anti-ransomware status of all volumes and storage VMs

You can view a list of all volumes and storage VMs (SVMs) and their current anti-ransomware status so you can determine whether the data on your volumes and SVMs are adequately protected from ransomware attacks.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

For more information on the different anti-ransomware statuses, see [ONTAP: Enable anti-ransomware](#).

View security details of all volumes with anti-ransomware detection

Steps

1. In the left navigation pane, click **Storage > Volumes**.
2. In the View menu, select **Health > Security > Anti-ransomware**
3. In the **Security: Anti-ransomware** view, you can sort by the various fields or use the Filter.



Anti-ransomware is not supported for offline volumes, restricted volumes, SnapLock volumes, FlexGroup volumes, FlexCache volumes, SAN-only volumes, volumes of stopped storage VMs, root volumes of storage VMs, or data protection volumes.

View security details of all storage VMs with anti-ransomware detection

Steps

1. In the left navigation pane, click **Storage > Storage VMs**.
2. Select **View > Security > Anti-ransomware**. A list of SVMs with the anti-ransomware status is displayed.



Anti-ransomware monitoring is not supported on storage VMs that do not have NAS protocol enabled.

Viewing all active security events

You can view all the active security events and then assign each of them to a user who can resolve the issue. Additionally, if there are certain security events that you do not want to receive, this list can help you identify the events that you want to disable.

What you'll need

You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Event Management**.

By default, New and Acknowledged events are displayed on the Event Management inventory page.

2. From the View menu, select **Active security events**.

The page displays all New and Acknowledged Security events that have been generated in the past 7 days.

Adding alerts for security events

You can configure alerts for individual security events just like any other events received by Unified Manager. Additionally, if you want to treat all security events alike and have email sent to the same person, you can create a single alert to notify you when any security events are triggered.

What you'll need

You must have the Application Administrator or Storage Administrator role.

The example below shows how to create an alert for the “Telnet Protocol Enabled” security event. This will send an alert if Telnet access is configured for remote administrative access into the cluster. You can use this same methodology to create alerts for all security events.

Steps

1. In the left navigation pane, click **Storage Management > Alert Setup**.
2. In the **Alert Setup** page, click **Add**.
3. In the **Add Alert** dialog box, click **Name**, and enter a name and description for the alert.
4. Click **Resources** and select the cluster or cluster on which you want to enable this alert.
5. Click **Events** and perform the following actions:
 - a. In the Event Severity list, select **Warning**.
 - b. In the Matching Events list, select **Telnet Protocol Enabled**.
6. Click **Actions** and then select the name of the user who will receive the alert email in the **Alert these users** field.
7. Configure any other options on this page for notification frequency, issuing SNMP traps, and executing a script.
8. Click **Save**.

Disabling specific security events

All events are enabled by default. You can disable specific events to prevent the generation of notifications for those events that are not important in your environment. You can enable events that are disabled if you want to resume receiving notifications for them.

What you'll need

You must have the Application Administrator or Storage Administrator role.

When you disable events, the previously generated events in the system are marked obsolete, and the alerts that are configured for these events are not triggered. When you enable events that are disabled, the notifications for these events are generated starting with the next monitoring cycle.

Steps

1. In the left navigation pane, click **Storage Management > Event Setup**.
2. In the **Event Setup** page, disable or enable events by choosing one of the following options:

If you want to...	Then do this...
Disable events	<ol style="list-style-type: none">a. Click Disable.b. In the Disable Events dialog box, select the Warning severity. This is the category for all security events.c. In the Matching Events column, select the security events that you want to disable, and then click the right arrow to move those events to the Disable Events column.d. Click Save and Close.e. Verify that the events that you disabled are displayed in the list view of the Event Setup page.
Enable events	<ol style="list-style-type: none">a. From the list of disabled events, select the check box for the event, or events, that you want to reenable.b. Click Enable.

Security events

Security events provide you with information about the security status of ONTAP clusters, storage virtual machines (SVMs), and volumes based on parameters defined in the *NetApp Security Hardening Guide for ONTAP 9*. These events notify you of potential issues so that you can evaluate their severity and fix the issue if necessary.

Security events are grouped by source type and include the event and trap name, impact level, and severity. These events appear in the cluster and storage VM event categories.

Managing backup and restore operations

You can create backups of Active IQ Unified Manager and use the restore feature to restore the backup to the same (local) system or a new (remote) system in case of a system failure or data loss.

There are three backup and restore methods depending on the operating system on which you have installed Unified Manager, and based on the number of clusters and nodes being managed:

Operating System	Size of Deployment	Recommended Backup Method
VMware vSphere	Any	VMware snapshot of the Unified Manager virtual appliance
Red Hat Enterprise Linux or CentOS Linux	Small	Unified Manager MySQL database dump
	Large	NetApp Snapshot of Unified Manager database
Microsoft Windows	Small	Unified Manager MySQL database dump
	Large	NetApp Snapshot of Unified Manager database with iSCSI protocol

These different methods are described in the sections that follow.

Backup and restore for Unified Manager on virtual appliance

The backup and restore model for Unified Manager when installed on a virtual appliance is to capture and restore an image of the full virtual application.

The following tasks enable you to complete a backup of the virtual appliance:

1. Power off the VM and take a VMware snapshot of the Unified Manager virtual appliance.
2. Make a NetApp Snapshot copy on the datastore to capture the VMware snapshot.

If the datastore is not hosted on a system running ONTAP software, follow the storage vendor guidelines to create a backup of the VMware snapshot.

3. Replicate the NetApp Snapshot copy, or snapshot equivalent, to alternate storage.
4. Delete the VMware snapshot.

You should implement a backup schedule using these tasks to ensure that the Unified Manager virtual appliance is protected if issues arise.

To restore the VM, you can use the VMware snapshot you created to restore the VM to the backup point-in-time state.

Backup and restore using a MySQL database dump


A MySQL database dump backup is a copy of the Active IQ Unified Manager database and configuration files that you can use in case of a system failure or data loss. You can schedule a backup to be written to a local destination or to a remote destination. It is highly recommended that you define a remote location that is external to the Active IQ Unified Manager host system.



MySQL database dump is the default backup mechanism when Unified Manager is installed on a Linux and Windows server. However, if Unified Manager is managing a large number of cluster and nodes, or if your MySQL backups are taking many hours to complete, you can back up using Snapshot copies. This functionality is available on Red Hat Enterprise Linux, CentOS Linux systems, and Windows.

A database dump backup consists of a single file in the backup directory and one or more files in the database repository directory. The file in the backup directory is very small because it contains only a pointer to the files located in the database repository directory that are required to recreate the backup.

The first time you generate a database backup a single file is created in the backup directory and a full backup file is created in the database repository directory. The next time you generate a backup a single file is created in the backup directory and an incremental backup file is created in the database repository directory that contains the differences from the full backup file. This process continues as you create additional backups, up to the maximum retention setting, as shown in the following figure.



Do not rename or remove any of the backup files in these two directories or any subsequent restore operation will fail.

If you write your backup files to the local system, you should initiate a process to copy the backup files to a remote location so they will be available in case you have a system issue that requires a complete restore.

Before beginning a backup operation, Active IQ Unified Manager performs an integrity check to verify that all the required backup files and backup directories exist and are writable. It also checks that there is enough space on the system to create the backup file.

Configuring the destination and schedule for database dump backups

You can configure the Unified Manager database dump backup settings to set the database backup path, retention count, and backup schedule. You can enable daily or weekly scheduled backups. By default, scheduled backups are disabled, but you should

set a backup schedule.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have a minimum of 150 GB of space available in the location you define as the backup path.

It is recommended that you use a remote location that is external to the Unified Manager host system.

- When Unified Manager is installed on a Linux system, and using MySQL backup, ensure that the following permissions and ownerships are set on the backup directory.

Permissions: 0750, Ownership: jboss:maintenance

- When Unified Manager is installed on a Windows system, and using MySQL backup, ensure that only the administrator has access to the backup directory.

More time is required the first time a backup is performed than for subsequent backups because the first backup is a full backup. A full backup can be over 1 GB and can take three to four hours. Subsequent backups are incremental and require less time.



- If you find the number of incremental backup files to be too large for the space you have allocated for backups, you can take a full backup periodically to replace the old backup and its incremental files. As another option, you can take a backup by using Snapshot copies.
- Backup taken during initial 15 days of a new cluster addition might not be accurate enough to get the historical performance data.

Steps

1. In the left navigation pane, click **General > Database Backup**.
2. In the **Database Backup** page, click **Backup Settings**.
3. Configure the appropriate values for a backup path, retention count, and schedule.

The default value for retention count is 10; you can use 0 for creating unlimited backups.

4. Select the **Scheduled Daily** or **Scheduled Weekly** button, and then specify the schedule details.
5. Click **Apply**.

Database dump backup files are created based on the schedule. You can see the available backup files in the Database Backup page.

What a database restore is

A MySQL database restore is the process of restoring an existing Unified Manager backup file to the same or a different Unified Manager server. You perform the restore operation from the Unified Manager maintenance console.

If you are performing a restore operation on the same (local) system, and the backup files are all stored locally, you can run the restore option using the default location. If you are performing a restore operation on a different Unified Manager system (a remote system), you must copy the backup file, or files, from secondary storage to the local disk before running the restore option.

During the restore process, you are logged out of Unified Manager. You can log in to the system after the

restore process is complete.

If you are restoring the backup image to a new server, after the restore operation completes you need to generate a new HTTPS security certificate and restart the Unified Manager server. You will also need to reconfigure SAML authentication settings, if they are required, when restoring the backup image to a new server.



Old backup files cannot be used to restore an image after Unified Manager has been upgraded to a newer version of software. To save space, all old backup files, except the newest file, are removed automatically when you upgrade Unified Manager.

Related information

[Generating an HTTPS security certificate](#)

[Enabling SAML authentication](#)

[Authentication with Active Directory or OpenLDAP](#)

Restoring a MySQL database backup on a Linux system

If data loss or data corruption occurs, you can restore Unified Manager to the previous stable state with minimum loss of data. You can restore the Unified Manager database to a local or remote Red Hat Enterprise Linux or CentOS system by using the Unified Manager maintenance console.

What you'll need

- You must have the root user credentials for the Linux host on which Unified Manager is installed.
- You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.
- You must have copied the Unified Manager backup file and the contents of the database repository directory to the system on which you will perform the restore operation.

It is recommended that you copy the backup file to the default directory `/data/ocum-backup`. The database repository files must be copied to the `/database-dumps-repo` subdirectory under the `/ocum-backup` directory.

- The backup files must be of `.7z` type.

The restore feature is platform-specific and version-specific. You can restore a Unified Manager backup only on the same version of Unified Manager. You can restore a Linux backup file or a virtual appliance backup file to a Red Hat Enterprise Linux or CentOS system.



If the backup folder name contains a space, you must include the absolute path or relative path in double quotation marks.

Steps

1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.

2. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system.
3. Log in to the system with the maintenance user (umadmin) name and password.
4. Enter the command `maintenance_console` and press Enter.
5. In the maintenance console **Main Menu**, enter the number for the **Backup Restore** option.
6. Enter the number for the **Restore MySQL Backup**.
7. When prompted, enter the absolute path of the backup file.

```
Bundle to restore from: /data/ocum-  
backup/UM_9.8.N151113.1348_backup_rhel_02-20-2020-04-45.7z
```

After the restore operation is complete, you can log in to Unified Manager.

After you restore the backup, if the OnCommand Workflow Automation server does not work, perform the following steps:

1. On the Workflow Automation server, change the IP address of the Unified Manager server to point to the latest machine.
2. On the Unified Manager server, reset the database password if the acquisition fails in step 1.

Restoring a MySQL database backup on Windows

In case of data loss or data corruption, you can use the restore feature to restore Unified Manager to the previous stable state with minimal loss. You can restore the Unified Manager MySQL database to a local Windows system or a remote Windows system by using the Unified Manager maintenance console.

What you'll need

- You must have Windows administrator privileges.
- You must have copied the Unified Manager backup file and the contents of the database repository directory to the system on which you will perform the restore operation.

It is recommended that you copy the backup file to the default directory `\ProgramData\NetApp\OnCommandAppData\ocum\backup`. The database repository files must be copied to the `\database_dumps_repo` subdirectory under the `\backup` directory.

- The backup files must be of `.7z` type.

The restore feature is platform-specific and version-specific. You can restore a Unified Manager MySQL backup only on the same version of Unified Manager, and a Windows backup can be restored only on a Windows platform.



If the folder names contain a space, you must include the absolute path or relative path of the backup file in double quotation marks.

Steps

1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
2. Log in to the Unified Manager system with administrator credentials.
3. Launch PowerShell or command prompt as a Windows administrator.
4. Enter the command `maintenance_console` and press Enter.
5. In the maintenance console **Main Menu**, enter the number for the **Backup Restore** option.
6. Enter the number for the **Restore MySQL Backup**.
7. When prompted, enter the absolute path of the backup file.

```
Bundle to restore from:  
\ProgramData\NetApp\OnCommandAppData\ocum\backup\UM_9.8.N151118.2300_bac  
kup_windows_02-20-2020-02-51.7z
```

After the restore operation is complete, you can log in to Unified Manager.

After you restore the backup, if the OnCommand Workflow Automation server does not work, perform the following steps:

1. On the Workflow Automation server, change the IP address of the Unified Manager server to point to the latest machine.
2. On the Unified Manager server, reset the database password if the acquisition fails in step 1.

Backup and restore using NetApp Snapshots

A NetApp Snapshot copy creates a point-in-time image of the Unified Manager database and configuration files that you can use to restore in case of a system failure or data loss. You schedule a Snapshot copy to be written to a volume on one of your ONTAP clusters periodically so that you always have a current copy.



This functionality is not available for Active IQ Unified Manager installed on a virtual appliance.

Configuring backup on Linux

If your Active IQ Unified Manager is installed on a Linux machine, then you can decide to configure your backup and restore using NetApp Snapshots.

Snapshot copies take very little time, usually just a few minutes, and the Unified Manager database is locked for a very short timeframe, so there is very little disruption to your installation. The image consumes minimal storage space and incurs negligible performance overhead because it records only changes to files since the last Snapshot copy was made. Because the Snapshot is created on an ONTAP cluster, you can take advantage of other NetApp features such as SnapMirror to create secondary protection, if needed.

Before beginning a backup operation, Unified Manager performs an integrity check to verify that the destination system is available.



- You can restore a Snapshot copy only on the same version of Active IQ Unified Manager.

For example, if you created a backup on Unified Manager 9.13, the backup can be restored only on Unified Manager 9.13 systems.

- If there is any change in the Snapshot configuration, it might cause the snapshot to be invalid.

Configuring Snapshot copy location

You can configure the volume where Snapshot copies will be stored on one of your ONTAP clusters using ONTAP System Manager or using the ONTAP CLI.

What you'll need

The cluster, storage VM, and volume must meet the following requirements:

- Cluster requirements:
 - ONTAP 9.3 or greater must be installed
 - It should be geographically close to the Unified Manager server
 - It can be monitored by Unified Manager, but it is not required
- Storage VM requirements:
 - The name switch and name mapping must be set to use “files”
 - Local users created to correspond with client-side users
 - Make sure All Read/Write access is selected
 - Make sure that Superuser Access is set to “any” in the export policy
 - NFS for NetApp Snapshot for Linux
 - NFSv4 must be enabled on the NFS server and NFSv4 ID domain specified on the client and storage VM
 - The volume should be at least double the size of the Unified Manager/opt/netapp/data directory

Use the command `du -sh /opt/netapp/data/` to check the current size.

- Volume requirements:
 - The volume should be at least double the size of the Unified Manager /opt/netapp/data directory
 - The security style must be set to UNIX
 - The local snapshot policy must be disabled
 - Volume autosize should be enabled
 - The performance service level should be set to a policy with high IOPS and low latency, such as “Extreme”

For detailed steps to create the NFS volume, see [How to configure NFSv4 in ONTAP 9](#) and the [ONTAP 9 NFS Configuration Express Guide](#).

Specifying the destination location for Snapshot copies

You should configure the destination location for Active IQ Unified Manager Snapshot copies on a volume that you have already configured in one of your ONTAP clusters. You should use maintenance console to define the location.

- You must have the root user credentials for the Linux host on which Active IQ Unified Manager is installed.
- You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.
- You must have the Cluster Management IP address, the name of the storage VM, the name of the volume, and the storage system user name and password.
- You must have mounted the volume to the Active IQ Unified Manager host, and you must have the mount path.

Steps

1. Use Secure Shell to connect to the IP address or FQDN of the Active IQ Unified Manager system.
2. Log in to the system with the maintenance user (umadmin) name and password.
3. Enter the command `maintenance_console` and press Enter.
4. In the maintenance console **Main Menu**, enter the number for the **Backup Restore** option.
5. Enter the number for **Configure NetApp Snapshot Backup**.
6. Enter the number to configure NFS.
7. Review the information that you will need to provide and then enter the number for **Enter Backup Configuration Details**.
8. To identify the volume where the Snapshot will be written, enter the IP address of the Cluster Management interface, the name of the storage VM, the name of the volume, LUN name, the storage system user name and password, and the mount path.
9. Verify this information and enter `y`.

The system performs the following tasks:

- Establishes the connection to the cluster
 - Stops all the services
 - Creates a new directory in the volume and copies the Active IQ Unified Manager database configuration files
 - Deletes the files from Active IQ Unified Manager and creates a symlink to the new database directory
 - Restarts all the services
10. Exit the maintenance console and launch the Active IQ Unified Manager interface to create a schedule for the Snapshot copy if you have not already done this.

Configuring backup on Windows

Active IQ Unified Manager supports backup and restore using NetApp Snapshots on Windows operating system with the help of LUN using iSCSI protocol.

Snapshot based backup can be taken while all Unified Manager services are running. A consistent state of database is captured as part of the Snapshot as the backup puts a global read lock on the entire database that

prevents any concurrent write. For your Unified Manager system installed on Windows OS to perform backup and restore using NetApp Snapshots, you should first configure Unified Manager backup to Snapshot based using the maintenance console.

Before you configure Unified Manager for creating Snapshot copies, you should perform the following configuration tasks.

- Configure ONTAP cluster
- Configure Windows host machine

Configuring backup location for Windows

You should configure the volume for storing Snapshot copies after backing up Unified Manager on Windows.

What you'll need

The cluster, storage VM, and volume must meet the following requirements:

- Cluster requirements:
 - ONTAP 9.3 or greater must be installed
 - It should be geographically close to the Unified Manager server
 - It is monitored by Unified Manager
- Storage VM requirements:
 - iSCSI connectivity on ONTAP cluster
 - iSCSI protocol must be enabled for the configured machine
 - You should have a dedicated volume and LUN for backup configuration. The selected volume should contain only one LUN and nothing else.
 - The size of the LUN should be at least twice the data size expected to be handled in the 9.9 Active IQ Unified Manager.

This sets the same size requirement on volume as well.
 - Make sure All Read/Write access is selected
 - Make sure that Superuser Access is set to “any” in the export policy
- Volume and LUN requirements:
 - The volume should be at least double the size of the Unified Manager MySQL data directory.
 - The security style must be set to Windows
 - The local snapshot policy must be disabled
 - Volume autosize should be enabled
 - The performance service level should be set to a policy with high IOPS and low latency, such as “Extreme”

Configuring ONTAP cluster

You need to perform few pre-configuration steps on ONTAP clusters before you can back up and restore Active IQ Unified Manager using Snapshot copy on Windows systems.

You can configure ONTAP cluster using either the command prompt or System Manager user interface. The configuration of ONTAP cluster involves configuring Data LIFs to be available to be assigned as iSCSI LIFs to the storage VM. The next step is to configure an iSCSI enabled storage VM using the System Manager user interface. You will need to configure a static network route for this storage VM to control how LIFs use the network for outbound traffic.



You should have a dedicated volume and a LUN for backup configuration. The selected volume should include only one LUN. The size of the LUN should be at least twice the data size expected to be handled by Active IQ Unified Manager.

You need to perform the following configuration:

Steps

1. Configure a iSCSI enabled storage VM or use an existing storage VM that has the same configuration.
2. Configure a network route for the configured storage VM.
3. Configure a volume of appropriate capacity and a single LUN inside ensuring that the volume is dedicated only for this LUN.



In a scenario when the LUN is created on System Manager, unmapping the LUN might delete the igroup, and restore might fail. To avoid this scenario, ensure that while creating a LUN, it is created explicitly and is not deleted when the LUN is unmapped.

4. Configure an initiator group in the storage VM.
5. Configure a port set.
6. Integrate the igroup with the portset.
7. Map the LUN to the igroup.

Configuring Windows host machine

You need to configure your Windows host machine before you can use NetApp Snapshot to back up and restore Active IQ Unified Manager. To start the Microsoft iSCSI initiator on a Windows host machine, type “iscsi” in the search bar and click **iSCSI Initiator**.

What you'll need

You should clean up any previous configurations on the host machine.

If you are trying to start the iSCSI initiator on a fresh installation of Windows, you are prompted for confirmation, and on your confirmation, the iSCSI Properties dialog box is displayed. If it is an existing Windows installation, then the iSCSI Properties dialog box displayed with a target that is either inactive or trying to connect. So, you will need to ensure that all the previous configurations on the Windows host are removed.

Steps

1. Clean up any previous configurations on the host machine.
2. Discover the target portal.
3. Connect to the target portal.
4. Connect using multipath to the target portal.

5. Discover both the LIFs.
6. Discover the LUN configured in the Windows machine as a device.
7. Configure the discovered LUN as a new volume drive in Windows.

Specifying the destination location for Snapshot copies on Windows

You should configure the destination location for Active IQ Unified Manager Snapshot copies on a volume that you have already configured in one of your ONTAP clusters. You should use maintenance console to define the location.

- You must have the administrator privilege for Windows host on which Active IQ Unified Manager is installed.
- You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.
- You must have the Cluster Management IP address, the name of the storage VM, the name of the volume, LUN name, and the storage system user name and password.
- You must have mounted the volume as a network drive to the Active IQ Unified Manager host, and you must have the mount drive.

Steps

1. Using Power Shell, connect to the IP address or fully qualified domain name of the Active IQ Unified Manager system.
2. Log in to the system with the maintenance user (umadmin) name and password.
3. Enter the command `maintenance_console` and press Enter.
4. In the maintenance console **Main Menu**, enter the number for the **Backup Restore** option.
5. Enter the number for **Configure NetApp Snapshot Backup**.
6. Enter the number to configure iSCSI.
7. Review the information that you will need to provide and then enter the number for **Enter Backup Configuration Details**.
8. To identify the volume where the Snapshot will be written, enter the IP address of the Cluster Management interface, the name of the storage VM, the name of the volume, LUN name, the storage system user name and password, and the mount drive.
9. Verify this information and enter `y`.

The system performs the following tasks:

- Storage VM is validated
- Volume is validated
- Mount drive and status is validated
- LUN existence and status
- Network drive existence
- Existence of recommend space (more than twice of mysql data directory) at mounted volume is validated
- LUN path corresponding to the dedicated LUN in the volume

- igroup name
- GUID of the volume where the network drive is mounted
- iSCSI initiator used to communicate with ONTAP

10. Exit the maintenance console and launch the Active IQ Unified Manager interface to create a schedule for Snapshot copies.

Configuring backup by Snapshot copy from maintenance console

To take Active IQ Unified Manager backup by using Snapshot copy, you should perform a few configuration steps from the maintenance console.

What you'll need

You should have the following details for your system:

- Cluster IP address
- Storage VM name
- Volume name
- LUN name
- Mount path
- Storage system credentials

Steps

1. Access the maintenance console of Unified Manager.
2. Enter 4 to select **Backup Restore**.
3. Enter 2 to select **Backup and Restore using NetApp Snapshot**.



If you want to change the backup configuration, then enter 3 for selecting **Update NetApp Snapshot Backup Configuration**. You can only update the password.

4. From the menu, enter 1 to select the **Configure NetApp Snapshot Backup**.
5. Enter 1 to provide the required information.
6. Provide the username and password for the maintenance console, and then provide the confirmation that LUN is mounted on host.

The process then verifies that the data directory, LUN path, storage VM, volumes, space availability, drive, and so on provided by you are correct. The operations that proceed in the background are:

- Services are stopped
- Database directory is moved to mounted storage
- Database directory is deleted and symlinks are established
- Services are restarted After the configuration completes in the Active IQ Unified Manager interface, the backup type is modified to NetApp Snapshot and reflects in the user interface as Database backup (Snapshot based).

Before beginning a backup operation, you must check whether there is any change in the Snapshot configuration because it might cause the snapshot to be invalid. Suppose you configured backup in G drive

and Snapshot taken. You later reconfigured the backup to E drive and data is saved to E drive as per the new configuration. If you try to restore Snapshot taken while it was in G drive, it fails with error that G drive does not exist.

Defining a backup schedule for Linux and Windows

You can configure the schedule at which Unified Manager Snapshot copies are created by using the Unified Manager UI.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have configured the settings for creating Snapshot copies from the maintenance console to identify the destination where the snapshots will be created.

Snapshot copies are created in just a few minutes and the Unified Manager database is locked only for few seconds.



Backup taken during initial 15 days of a new cluster addition might not be accurate enough to get the historical performance data.

Steps

1. In the left navigation pane, click **General > Database Backup**.
2. In the **Database Backup** page, click **Backup Settings**.
3. Enter the maximum number of Snapshot copies that you want to retain in the **Retention Count** field.

The default value for retention count is 10. The maximum number of Snapshot copies is determined by the version of ONTAP software on the cluster. You can leave this field blank to implement the maximum value regardless of ONTAP version.

4. Select the **Scheduled Daily** or **Scheduled Weekly** button, and then specify the schedule details.
5. Click **Apply**.

Snapshot copies are created based on the schedule. You can see the available backup files in the Database Backup page.

Because of the importance of this volume and the snapshots, you may want to create one or two alerts for this volume so you are notified when either:

- The volume space is 90% full. Use the event **Volume Space Full** to set up the alert.

You can add capacity to the volume using ONTAP System Manager or the ONTAP CLI so that the Unified Manager database does not run out of space.

- The number of snapshots is close to reaching the maximum number. Use the event **Too Many Snapshot Copies** to set up the alert.

You can delete older snapshots using ONTAP System Manager or the ONTAP CLI so that there is always room for new Snapshot copies.

You configure alerts in the Alert Setup page.

Restoring Unified Manager by using Snapshot copies

If data loss or data corruption occurs, you can restore Unified Manager to the previous stable state with minimum loss of data. You can restore the Unified Manager Snapshot database to a local or remote operating system by using the Unified Manager maintenance console.

What you'll need

- You must have the root user credentials for the Linux host and administrative privileges for Windows host machine on which Unified Manager is installed.
- You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.

The restore feature is platform-specific and version-specific. You can restore a Unified Manager backup only on the same version of Unified Manager.

Steps

1. Connect to the IP address or fully qualified domain name of the Unified Manager system.
 - Linux: Secure Shell
 - Windows: Power Shell
2. Log in to the system with the root user credentials.
3. Enter the command `maintenance_console` and press Enter.
4. In the maintenance console **Main Menu**, enter 4 for the **Backup Restore** option.
5. Enter 2 for selecting **Backup and Restore using NetApp Snapshot**.

If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. Enter 1 for selecting **Configure NetApp Snapshot Backup** and configure the settings for Snapshot copies as they are on the original system.

6. Enter 3 for selecting **Restore using NetApp Snapshot**.
7. Select the Snapshot copy from which you want to restore Unified Manager. Press **Enter**.
8. After the restore process is complete, log in to the Unified Manager user interface.

After you restore the backup, if the Workflow Automation server does not work, perform the following steps:

1. On the Workflow Automation server, change the IP address of the Unified Manager server to point to the latest machine.
2. On the Unified Manager server, reset the database password if the acquisition fails in step 1.

Modifying the backup type

If you want to change the type of backup for your Active IQ Unified Manager system, then you can use the maintenance console options. The **Unconfigure NetApp Snapshot Backup** option enables you to fall back to the MySQL based backup.

What you'll need

You must have a user ID and password authorized to log in to the maintenance console of the Unified Manager server.

Steps

1. Access the maintenance console.
2. Select 4 from the **Main Menu** for backup and restore.
3. Select 2 from the **Backup and Restore Menu**.
4. Select 4 for **Unconfigure NetApp Snapshot Backup**.

The actions that are performed are displayed which are, stop the services, break the symlink, move the data from storage to directory, and then start the services again.

After the backup method is modified, the backup mechanism is changed from Snapshot copy to default MySQL backup. This change appears in the Database Backup section of the General settings.

On-demand backup for Unified Manager

You can use the Active IQ Unified Manager user interface to generate on demand backup whenever required. The on-demand backup enables you to instantaneously create a backup using the existing backup method. The on-demand backup does not differentiate between MySQL or NetApp Snapshot based backup.

You can perform on-demand backup using the **Backup Now** button on the Database Backup page. The on-demand backup does not depend on the schedules that you have configured for Active IQ Unified Manager.

Migrating a Unified Manager virtual appliance to a Linux system

You can restore a Unified Manager MySQL database dump backup from a virtual appliance to a Red Hat Enterprise Linux or CentOS Linux system if you want to change the host operating system on which Unified Manager is running.

What you'll need

- On the virtual appliance:
 - You must have the Operator, Application Administrator, or Storage Administrator role.
 - You must know the name of the Unified Manager maintenance user for the restore operation.
- On the Linux system:
 - You must have installed Unified Manager on a Linux server following the instructions in [Installing Unified Manager on Linux systems](#).
 - The version of Unified Manager on this server must be the same as the version on the virtual appliance from which you are using the backup file.
 - Do not launch the UI or configure any clusters, users, or authentication settings on the Linux system after installation. The backup file populates this information during the restore process.
 - You must have the root user credentials for the Linux host.

These steps describe how to create a backup file on the virtual appliance, copy the backup files to the Red Hat Enterprise Linux or CentOS system, and then restore the database backup to the new system.

Steps

1. On the virtual appliance, click **Management > Database Backup**.
2. In the **Database Backup** page, click **Backup Settings**.
3. Change the backup path to `/jail/support`.
4. In the Schedule section, select **Scheduled Daily**, and enter a time a few minutes past the current time so that the backup is created shortly.
5. Click **Apply**.
6. Wait a few hours for the backup to be generated.

A full backup can be over 1 GB and can take three to four hours to complete.

7. Log in as the root user to the Linux host on which Unified Manager is installed and copy the backup files from `/support` on the virtual appliance using SCP.

```
root@<rhel_server>:/# scp -r admin@<vapp_server_ip_address>:/support/* .
```

```
root@ocum_rhel-21:/# scp -r admin@10.10.10.10:/support/* .
```

Make sure you have copied the .7z backup file and all the .7z repository files in the `/database-dumps-repo` subdirectory.

8. At the command prompt, restore the backup:

```
um backup restore -f /<backup_file_path>/<backup_file_name>
```

```
um backup restore -f /UM_9.7.N151113.1348_backup_unix_02-12-2019-04-16.7z
```

9. After the restore operation completes, log in to the Unified Manager web UI.

You should perform the following tasks:

- Generate a new HTTPS security certificate and restart the Unified Manager server.
- Change the backup path to the default setting for your Linux system (`/data/ocum-backup`), or to a new path of your choice, because there is no `/jail/support` path on the Linux system.
- Reconfigure both sides of your Workflow Automation connection, if WFA is being used.
- Reconfigure SAML authentication settings, if you are using SAML.

After you have verified that everything is running as expected on your Linux system, you can shut down and remove the Unified Manager virtual appliance.

Managing scripts

You can use scripts to automatically modify or update multiple storage objects in Unified Manager. The script is associated with an alert. When an event triggers an alert, the script is executed. You can upload custom scripts and test their execution when an alert is generated.

The ability to upload scripts to Unified Manager and run them is enabled by default. If your organization does not want to allow this functionality because of security reasons, you can disable this functionality from **Storage Management > Feature Settings**.

How scripts work with alerts

You can associate an alert with your script so that the script is executed when an alert is raised for an event in Unified Manager. You can use the scripts to resolve issues with storage objects or identify which storage objects are generating the events.

When an alert is generated for an event in Unified Manager, an alert email is sent to the specified recipients. If you have associated an alert with a script, the script is executed. You can get the details of the arguments passed to the script from the alert email.



If you have created a custom script and associated it with an alert for a specific event type, actions are taken based on your custom script for that event type, and the **Fix it** actions are not available by default on the Management Actions page or Unified Manager dashboard.

The script uses the following arguments for execution:

- -eventID
- -eventName
- -eventSeverity
- -eventSourceID
- -eventSourceName
- -eventSourceType
- -eventState
- -eventArgs

You can use the arguments in your scripts and gather related event information or modify storage objects.

Example for obtaining arguments from scripts

```
print "$ARGV[0] : $ARGV[1]\n"  
print "$ARGV[7] : $ARGV[8]\n"
```

When an alert is generated, this script is executed and the following output is displayed:

```
-eventID : 290  
-eventSourceID : 4138
```

Adding scripts

You can add scripts in Unified Manager, and associate the scripts with alerts. These scripts are executed automatically when an alert is generated, and enable you to obtain information about storage objects for which the event is generated.

What you'll need

- You must have created and saved the scripts that you want to add to the Unified Manager server.
- The supported file formats for scripts are Perl, Shell, PowerShell, Python, and .bat files.

Platform on which Unified Manager is installed	Supported languages
VMware	Perl and Shell scripts
Linux	Perl, Python, and Shell scripts
Windows	PowerShell, Perl, Python, and .bat scripts

- For Perl scripts, Perl must be installed on the Unified Manager server. For VMware installations, Perl 5 is installed by default and scripts will support only what Perl 5 supports. If Perl was installed after Unified Manager, you must restart the Unified Manager server.
- For PowerShell scripts, the appropriate PowerShell execution policy must be set on the Windows server so that the scripts can be executed.



If your script creates log files to track the alert script progress, you must make sure that the log files are not created anywhere within the Unified Manager installation folder.

- You must have the Application Administrator or Storage Administrator role.

You can upload custom scripts and gather event details about the alert.



If you do not see this capability available in the user interface it is because the functionality has been disabled by your administrator. If required, you can enable this functionality from **Storage Management > Feature Settings**.

Steps

1. In the left navigation pane, click **Storage Management > Scripts**.
2. In the **Scripts** page, click **Add**.
3. In the **Add Script** dialog box, click **Browse** to select your script file.
4. Enter a description for the script that you select.
5. Click **Add**.

Deleting scripts

You can delete a script from Unified Manager when the script is no longer required or valid.

What you'll need

- You must have the Application Administrator or Storage Administrator role.
- The script must not be associated with an alert.

Steps

1. In the left navigation pane, click **Storage Management > Scripts**.

2. In the **Scripts** page, select the script that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Testing script execution

You can verify that your script is executed correctly when an alert is generated for a storage object.

- You must have the Application Administrator or Storage Administrator role.
- You must have uploaded a script in the supported file format to Unified Manager.

Steps

1. In the left navigation pane, click **Storage Management > Scripts**.
2. In the Scripts page, add your test script.
3. In the left navigation pane, click **Storage Management > Alert Setup**.
4. In the **Alert Setup** page, perform one of the following actions:

To...	Do this...
Add an alert	<ol style="list-style-type: none"> a. Click Add. b. In the Actions section, associate the alert with your test script.
Edit an alert	<ol style="list-style-type: none"> a. Select an alert, and then click Edit. b. In the Actions section, associate the alert with your test script.

5. Click **Save**.
6. In the **Alert Setup** page, select the alert that you added or modified, and then click **Test**.

The script is executed with the “-test” argument, and a notification alert is sent to the email addresses that were specified when the alert was created.

Managing and monitoring groups

You can create groups in Unified Manager to manage storage objects.

Understanding groups

You can create groups in Unified Manager to manage storage objects. Understanding the concepts about groups and how group rules enable you to add storage objects to a group will help you to manage the storage objects in your environment.

What a group is

A group is a dynamic collection of heterogeneous storage objects (clusters, SVMs, or volumes). You can create groups in Unified Manager to easily manage a set of storage

objects. The members in a group might change, depending on the storage objects that are monitored by Unified Manager at a point in time.

- Each group has a unique name.
- You must configure a minimum of one group rule for each group.
- You can associate a group with more than one group rule.
- Each group can include multiple types of storage objects such as clusters, SVMs, or volumes.
- Storage objects are dynamically added to a group based on when a group rule is created or when Unified Manager completes a monitoring cycle.
- You can simultaneously apply actions on all the storage objects in a group such as setting thresholds for volumes.

How group rules work for groups

A group rule is a criterion that you define to enable storage objects (volumes, clusters, or SVMs) to be included in a specific group. You can use condition groups or conditions for defining group rule for a group.

- You must associate a group rule to a group.
- You must associate an object type for a group rule; only one object type is associated for a group rule.
- Storage objects are added or removed from the group after each monitoring cycle or when a rule is created, edited, or deleted.
- A group rule can have one or more condition groups, and each condition group can have one or more conditions.
- Storage objects can belong to multiple groups based on group rules you create.

Conditions

You can create multiple condition groups, and each condition group can have one or more conditions. You can apply all the defined condition groups in a group rule for groups in order to specify which storage objects are included in the group.

Conditions within a condition group are executed using logical AND. All the conditions in a condition group must be met. When you create or modify a group rule, a condition is created that applies, selects, and groups only those storage objects that satisfy all conditions in the condition group. You can use multiple conditions within a condition group when you want to narrow the scope of which storage objects to include in a group.

You can create conditions with storage objects by using the following operands and operator and specifying the required value.

Storage object type	Applicable operands
Volume	<ul style="list-style-type: none">• Object name• Owning cluster name• Owning SVM name• Annotations

Storage object type	Applicable operands
SVM	<ul style="list-style-type: none"> • Object name • Owning cluster name • Annotations
Cluster	<ul style="list-style-type: none"> • Object name • Annotations

When you select annotation as an operand for any storage object, the “Is” operator is available. For all other operands, you can select either “Is” or “Contains” as operator.

- Operand

The list of operands in Unified Manager changes based on the selected object type. The list includes the object name, owning cluster name, owning SVM name, and annotations that you define in Unified Manager.

- Operator

The list of operators changes based on the selected operand for a condition. The operators supported in Unified Manager are “Is” and “Contains”.

When you select the “Is” operator, the condition is evaluated for exact match of operand value to the value provided for the selected operand.

When you select the “Contains” operator, the condition is evaluated to meet one of the following criteria:

- The operand value is an exact match to the value provided for the selected operand
- The operand value contains the value provided for the selected operand

- Value

The value field changes based on the operand selected.

Example of a group rule with conditions

Consider a condition group for a volume with the following two conditions:

- Name contains “vol”
- SVM name is “data_svm”

This condition group selects all volumes that include “vol” in their names and that are hosted on SVMs with the name “data_svm”.

Condition groups

Condition groups are executed using logical OR, and then applied to storage objects. The storage objects must satisfy one of the condition groups to be included in a group. The storage objects of all the condition groups are combined. You can use condition groups to increase the scope of storage objects to include in a group.

Example of a group rule with condition groups

Consider two condition groups for a volume, with each group containing the following two conditions:

- Condition group 1
 - Name contains “vol”
 - SVM name is “data_svm” Condition group 1 selects all volumes that include “vol” in their names and that are hosted on SVMs with the name “data_svm”.
- Condition group 2
 - Name contains “vol”
 - The annotation value of data-priority is “critical” Condition group 2 selects all volumes that include “vol” in their names and that are annotated with the data-priority annotation value as “critical”.

When a group rule containing these two condition groups is applied on storage objects, then the following storage objects are added to a selected group:

- All volumes that include “vol” in their names and that are hosted on the SVM with the name “data_svm”.
- All volumes that include “vol” in their names and that are annotated with the data-priority annotation value “critical”.

How group actions work on storage objects

A group action is an operation that is performed on all the storage objects in a group. For example, you can configure volume threshold group action to simultaneously change the volume threshold values of all volumes in a group.

Groups support unique group action types. You can have a group with only one volume health threshold group action type. However, you can configure a different type of group action, if available, for the same group. The rank of a group action determines the order in which the action is applied to storage objects. The details page of a storage object provides information about which group action is applied on the storage object.

Example of unique group actions

Consider a volume A that belongs to groups G1 and G2, and the following volume health threshold group actions are configured for these groups:

- `Change_capacity_threshold` group action with rank 1, for configuring the capacity of the volume
- `Change_snapshot_copies` group action with rank 2, for configuring the Snapshot copies of the volume

The `Change_capacity_threshold` group action always takes priority over the `Change_snapshot_copies` group action and is applied to volume A. When Unified Manager completes one cycle of monitoring, the health threshold related events of volume A are re-evaluated per the `Change_capacity_threshold` group action. You cannot configure another volume threshold type of group action for either G1 or G2 group.

Adding groups

You can create groups to combine clusters, volumes, and storage virtual machines (SVMs) for ease of management.

What you'll need

You must have the Application Administrator or Storage Administrator role.

You can define group rules to add or remove members from the group and to modify group actions for the group.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Groups** tab, click **Add**.
3. In the **Add Group** dialog box, enter a name and description for the group.
4. Click **Add**.

Editing groups

You can edit the name and description of a group that you created in Unified Manager.

What you'll need

You must have the Application Administrator or Storage Administrator role.

When you edit a group to update the name, you must specify a unique name; you cannot use an existing group name.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Groups** tab, select the group that you want to edit, and then click **Edit**.
3. In the **Edit Group** dialog box, change the name, description, or both for the group.
4. Click **Save**.

Deleting groups

You can delete a group from Unified Manager when the group is no longer required.

What you'll need

- None of the storage objects (clusters, SVMs, or volumes) must be associated with any group rule that is associated with the group that you want to delete.
- You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Groups** tab, select the group that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Deleting a group does not delete the group actions that are associated with the group. However, these group actions will be unmapped after the group is deleted.

Adding group rules

You can create group rules for a group to dynamically add storage objects such as volumes, clusters, or storage virtual machines (SVMs) to the group. You must configure at least one condition group with at least one condition to create a group rule.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Storage objects that are currently monitored are added as soon as the group rule is created. New objects are added only after the monitoring cycle is completed.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Rules** tab, click **Add**.
3. In the **Add Group Rule** dialog box, specify a name for the group rule.
4. In the **Target Object Type** field, select the type of storage object that you want to group.
5. In the **Group** field, select the required group for which you want to create group rules.
6. In the **Conditions** section, perform the following steps to create a condition, a condition group, or both:

To create....	Do this...
A condition	<ol style="list-style-type: none">a. Select an operand from the list of operands.b. Select either Contains or Is as the operator.c. Enter a value, or select a value from the available list.
A condition group	<ol style="list-style-type: none">a. Click Add Condition Groupb. Select an operand from the list of operands.c. Select either Contains or Is as the operator.d. Enter a value, or select a value from the available list.e. Click Add condition to create more conditions if required, and repeat steps a through d for each condition.

7. Click **Add**.

Example for creating a group rule

Perform the following steps in the Add Group Rule dialog box to create a group rule, including configuring a condition and adding a condition group:

Steps

1. Specify a name for the group rule.
2. Select the object type as storage virtual machine (SVM).

3. Select a group from the list of groups.
4. In the Conditions section, select **Object Name** as the operand.
5. Select **Contains** as the operator.
6. Enter the value as `svm_data`.
7. Click **Add condition group**.
8. Select **Object Name** as the operand.
9. Select **Contains** as the operator.
10. Enter the value as `vol`.
11. Click **Add condition**.
12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **critical** as the value in step 10.
13. Click **Add** to create the condition for the group rule.

Editing group rules

You can edit group rules to modify the condition groups and the conditions within a condition group to add or remove storage objects to or from a specific group.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Rules** tab, select the group rule that you want to edit, and then click **Edit**.
3. In the **Edit Group Rule** dialog box, change the group rule name, associated group name, condition groups, and conditions as required.



You cannot change the target object type for a group rule.

4. Click **Save**.

Deleting group rules

You can delete a group rule from Active IQ Unified Manager when the group rule is no longer required.

What you'll need

You must have the Application Administrator or Storage Administrator role.

When a group rule is deleted, the associated storage objects will be removed from the group.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Rules** tab, select the group rule that you want to delete, and then click **Delete**.

3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Adding group actions

You can configure group actions that you want to apply to storage objects in a group. Configuring actions for a group enables you to save time, because you do not have to add these actions to each object individually.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Actions** tab, click **Add**.
3. In the **Add Group Action** dialog box, enter a name and description for the action.
4. From the **Group** menu, select a group for which you want to configure the action.
5. From the **Action Type** menu, select an action type.

The dialog box expands, enabling you to configure the selected action type with required parameters.

6. Enter appropriate values for the required parameters to configure a group action.
7. Click **Add**.

Editing group actions

You can edit the group action parameters that you configured in Unified Manager, such as the group action name, description, associated group name, and parameters of the action type.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Actions** tab, select the group action that you want to edit, and then click **Edit**.
3. In the **Edit Group Action** dialog box, change the group action name, description, associated group name, and parameters of the action type, as required.
4. Click **Save**.

Configuring volume health thresholds for groups

You can configure group-level volume health thresholds for capacity, Snapshot copies, qtree quotas, growth, and inodes.

What you'll need

You must have the Application Administrator or Storage Administrator role.

The volume health threshold type of group action is applied only on volumes of a group.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Actions** tab, click **Add**.
3. Enter a name and description for the group action.
4. From the **Group** drop-down box, select a group for which you want to configure group action.
5. Select **Action Type** as the volume health threshold.
6. Select the category for which you want to set the threshold.
7. Enter the required values for the health threshold.
8. Click **Add**.

Deleting group actions

You can delete a group action from Unified Manager when the group action is no longer required.

What you'll need

You must have the Application Administrator or Storage Administrator role.

When you delete the group action for the volume health threshold, global thresholds are applied to the storage objects in that group. Any object-level health thresholds that are set on the storage object are not impacted.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Actions** tab, select the group action that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, confirm the deletion by clicking **Yes**.

Reordering group actions

You can change the order of the group actions that are to be applied to the storage objects in a group. Group actions are applied to storage objects sequentially based on their rank. The lowest rank is assigned to the group action that you configured last. You can change the rank of the group action depending on your requirements.

What you'll need

You must have the Application Administrator or Storage Administrator role.

You can select either a single row or multiple rows, and then perform multiple drag-and-drop operations to change the rank of group actions. However, you must save the changes for the re-prioritization to be reflected in the group actions grid.

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Actions** tab, click **Reorder**.

3. In the **Reorder Group Actions** dialog box, drag and drop the rows to rearrange the sequence of group actions as required.
4. Click **Save**.

Prioritizing storage object events using annotations

You can create and apply annotation rules to storage objects so that you can identify and filter those objects based on the type of annotation applied and its priority.

Understanding more about annotations

Understanding the concepts about annotations helps you to manage the events related to the storage objects in your environment.

What annotations are

An annotation is a text string (the name) that is assigned to another text string (the value). Each annotation name-value pair can be dynamically associated with storage objects using annotation rules. When you associate storage objects with predefined annotations, you can filter and view the events that are related to them. You can apply annotations to clusters, volumes, and storage virtual machines (SVMs).

Each annotation name can have multiple values; each name-value pair can be associated with a storage object through rules.

For example, you can create an annotation named “data-center” with the values “Boston” and “Canada”. You can then apply the annotation “data-center” with the value “Boston” to volume v1. When an alert is generated for any event on a volume v1 that is annotated with “data-center”, the generated email indicates the location of the volume, “Boston”, and this enables you to prioritize and resolve the issue.

How annotation rules work in Unified Manager

An annotation rule is a criterion that you define to annotate storage objects (volumes, clusters, or storage virtual machines (SVMs)). You can use either condition groups or conditions for defining annotation rules.

- You must associate an annotation rule to an annotation.
- You must associate an object type for an annotation rule; only one object type can be associated for an annotation rule.
- Unified Manager adds or removes annotations from storage objects after each monitoring cycle or when a rule is created, edited, deleted, or reordered.
- An annotation rule can have one or more condition groups, and each condition group can have one or more conditions.
- Storage objects can have multiple annotations. An annotation rule for a particular annotation can also use different annotations in the rule conditions to add another annotation to already annotated objects.

Conditions

You can create multiple condition groups, and each condition group can have one or more conditions. You can

apply all the defined condition groups in an annotation rule of an annotation in order to annotate storage objects.

Conditions within a condition group are executed using logical AND. All the conditions in a condition group must be met. When you create or modify an annotation rule, a condition is created that applies, selects, and annotates only those storage objects that meet all the conditions in the condition group. You can use multiple conditions within a condition group when you want to narrow the scope of which storage objects to annotate.

You can create conditions with storage objects by using the following operands and operator and specifying the required value.

Storage object type	Applicable operands
Volume	<ul style="list-style-type: none">• Object name• Owning cluster name• Owning SVM name• Annotations
SVM	<ul style="list-style-type: none">• Object name• Owning cluster name• Annotations
Cluster	<ul style="list-style-type: none">• Object name• Annotations

When you select annotation as an operand for any storage object, the “Is” operator is available. For all other operands, you can select either “Is” or “Contains” as operator. When you select the “Is” operator, the condition is evaluated for an exact match of the operand value with the value provided for the selected operand. When you select the “Contains” operator, the condition is evaluated to meet one of the following criteria:

- The operand value is an exact match to the value of the selected operand.
- The operand value contains the value provided for the selected operand.

Example of an annotation rule with conditions

Consider an annotation rule with one condition group for a volume with the following two conditions:

- Name contains “vol”
- SVM name is “data_svm”

This annotation rule annotates all volumes that include “vol” in their names and that are hosted on SVMs with the name “data_svm” with the selected annotation and the annotation type.

Condition groups

Condition groups are executed using logical OR, and then applied to storage objects. The storage objects must meet the requirements of one of the condition groups to be annotated. The storage objects that meet the conditions of all the condition groups are annotated. You can use condition groups to increase the scope of storage objects to be annotated.

Example of an annotation rule with condition groups

Consider an annotation rule with two condition groups for a volume; each group contains the following two conditions:

- Condition group 1
 - Name contains “vol”
 - SVM name is “data_svm” This condition group annotates all volumes that include “vol” in their names and that are hosted on SVMs with the name “data_svm”.
- Condition group 2
 - Name contains “vol”
 - The annotation value of data-priority is “critical” This condition group annotates all volumes that include “vol” in their names and that are annotated with the data-priority annotation value as “critical”.

When an annotation rule containing these two condition groups is applied on storage objects, then the following storage objects are annotated:

- All volumes that include “vol” in their names and that are hosted on SVM with the name “data_svm”.
- All volumes that include “vol” in their names and that are annotated with the data-priority annotation value as “critical”.

Description of predefined annotation values

Data-priority is a predefined annotation that has the values Mission critical, High, and Low. These values enable you to annotate storage objects based on the priority of data that they contain. You cannot edit or delete the predefined annotation values.

- **Data-priority:Mission critical**

This annotation is applied to storage objects that contain mission-critical data. For example, objects that contain production applications can be considered as mission critical.

- **Data-priority:High**

This annotation is applied to storage objects that contain high-priority data. For example, objects that are hosting business applications can be considered high priority.

- **Data-priority:Low**

This annotation is applied to storage objects that contain low-priority data. For example, objects that are on secondary storage, such as backup and mirror destinations, might be of low priority.

Adding annotations dynamically

When you create custom annotations, Unified Manager dynamically associates clusters, storage virtual machines (SVMs), and volumes with the annotations by using rules. These rules automatically assign the annotations to storage objects.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotations** page, click **Add Annotation**.
3. In the **Add Annotation** dialog box, type a name and description for the annotation.
4. Optional: In the **Annotation Values** section, click **Add** to add values to the annotation.
5. Click **Save**.

Adding values to annotations

You can add values to annotations, and then associate storage objects with a particular annotation name-value pair. Adding values to annotations helps you to manage storage objects more effectively.

What you'll need

You must have the Application Administrator or Storage Administrator role.

You cannot add values to predefined annotations.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotations** page, select the annotation to which you want to add a value and then click **Add** in the **Values** section.
3. In the **Add Annotation Value** dialog box, specify a value for the annotation.

The value that you specify must be unique for the selected annotation.
4. Click **Add**.

Deleting annotations

You can delete custom annotations and their values when they are no longer required.

What you'll need

- You must have the Application Administrator or Storage Administrator role.
- The annotation values must not be used in other annotations or group rules.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotations** tab, select the annotation that you want to delete.

The details of the selected annotation are displayed.
3. Click **Actions > Delete** to delete the selected annotation and its value.
4. In the warning dialog box, click **Yes** to confirm the deletion.

Viewing the annotation list and details

You can view the list of annotations that are dynamically associated with clusters, volumes, and storage virtual machines (SVMs). You can also view details such as the description, created by, created date, values, rules, and the objects associated with the annotation.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotations** tab, click the annotation name to view the associated details.

Deleting values from annotations

You can delete values associated with custom annotations when that value no longer applies to the annotation.

What you'll need

- You must have the Application Administrator or Storage Administrator role.
- The annotation value must not be associated with any annotation rules or group rules.

You cannot delete values from predefined annotations.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the annotations list in the **Annotations** tab, select the annotation from which you want to delete a value.
3. In the **Values** area of the **Annotations** tab, select the value you want to delete, and then click **Delete**.
4. In the **Warning** dialog box, click **Yes**.

The value is deleted and no longer displayed in the list of values for the selected annotation.

Creating annotation rules

You can create annotation rules that Unified Manager uses to dynamically annotate storage objects such as volumes, clusters, or storage virtual machines (SVMs).

What you'll need

You must have the Application Administrator or Storage Administrator role.

Storage objects that are currently monitored are annotated as soon as the annotation rule is created. New objects are annotated only after the monitoring cycle is completed.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotation Rules** tab, click **Add**.
3. In the **Add Annotation Rule** dialog box, specify a name for the annotation rule.
4. In the **Target Object Type** field, select the type of storage object that you want to annotate.

5. In the **Apply Annotation** fields, select the annotation and annotation value that you want to use.
6. In the Conditions section, perform the appropriate action to create a condition, a condition group, or both:

To create...	Do this...
A condition	<ol style="list-style-type: none"> a. Select an operand from the list of operands. b. Select either Contains or Is as the operator. c. Enter a value, or select a value from the available list.
A condition group	<ol style="list-style-type: none"> a. Click Add Condition Group. b. Select an operand from the list of operands. c. Select either Contains or Is as the operator. d. Enter a value, or select a value from the available list. e. Click Add condition to create more conditions if required, and repeat steps a through d for each condition.

7. Click **Add**.

Example of creating an annotation rule

Perform the following steps in the Add Annotation Rule dialog box to create an annotation rule, including configuring a condition and adding a condition group:

Steps

1. Specify a name for the annotation rule.
2. Select the target object type as storage virtual machine (SVM).
3. Select an annotation from the list of annotations, and specify a value.
4. In the Conditions section, select **Object Name** as the operand.
5. Select **Contains** as the operator.
6. Enter the value as `svm_data`.
7. Click **Add condition group**.
8. Select **Object Name** as the operand.
9. Select **Contains** as the operator.
10. Enter the value as `vol`.
11. Click **Add condition**.
12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **mission-critical** as the value in step 10.
13. Click **Add**.

Adding annotations manually to individual storage objects

You can manually annotate selected volumes, clusters, and SVMs without using annotation rules. You can annotate a single storage object or multiple storage objects, and specify the required name-value pair combination for the annotation.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Steps

1. Navigate to the storage objects you want to annotate:

To add annotation to...	Do this...
Clusters	<ol style="list-style-type: none">a. Click Storage > Clusters.b. Select one or more clusters.
Volumes	<ol style="list-style-type: none">a. Click Storage > Volumes.b. Select one or more volumes.
SVMs	<ol style="list-style-type: none">a. Click Storage > SVMs.b. Select one or more SVMs.

2. Click **Annotate** and select a name-value pair.
3. Click **Apply**.

Editing annotation rules

You can edit annotation rules to modify the condition groups and conditions within the condition group to add annotations to or remove annotations from storage objects.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Annotations are dissociated from storage objects when you edit the associated annotation rules.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotation Rules** tab, select the annotation rule you want to edit, and then click **Actions > Edit**.
3. In the **Edit Annotation Rule** dialog box, change the rule name, annotation name and value, condition groups, and conditions as required.

You cannot change the target object type for an annotation rule.

4. Click **Save**.

Configuring conditions for annotation rules

You can configure one or more conditions to create annotation rules that Unified Manager applies on the storage objects. The storage objects that satisfy the annotation rule are annotated with the value specified in the rule.

What you'll need

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotation Rules** tab, click **Add**.
3. In the **Add Annotation Rule** dialog box, enter a name for the rule.
4. Select one object type from the Target Object Type list, and then select an annotation name and value from the list.
5. In the **Conditions** section of the dialog box, select an operand and an operator from the list and enter a condition value, or click **Add Condition** to create a new condition.
6. Click **Save and Add**.

Example of configuring a condition for an annotation rule

Consider a condition for the object type SVM, where the object name contains "svm_data".

Perform the following steps in the Add Annotation Rule dialog box to configure the condition:

Steps

1. Enter a name for the annotation rule.
2. Select the target object type as SVM.
3. Select an annotation from the list of annotations and a value.
4. In the **Conditions** field, select **Object Name** as the operand.
5. Select **Contains** as the operator.
6. Enter the value as `svm_data`.
7. Click **Add**.

Deleting annotation rules

You can delete annotation rules from Active IQ Unified Manager when the rules are no longer required.

What you'll need

You must have the Application Administrator or Storage Administrator role.

When you delete an annotation rule, the annotation is disassociated and removed from the storage objects.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.

2. In the **Annotation Rules** tab, select the annotation rule that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, click **Yes** to confirm the deletion.

Reordering annotation rules

You can change the order in which Unified Manager applies annotation rules to storage objects. Annotation rules are applied to storage objects sequentially based on their rank. When you configure an annotation rule, the rank is least. But you can change the rank of the annotation rule depending on your requirements.

What you'll need

You must have the Application Administrator or Storage Administrator role.

You can select either a single row or multiple rows and perform many drag-and-drop operations to change the rank of annotation rules. However, you must save the changes for the reprioritization to be displayed in the Annotation Rules tab.

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotation Rules** tab, click **Reorder**.
3. In the **Reorder Annotation Rule** dialog box, drag and drop single or multiple rows to rearrange the sequence of the annotation rules.
4. Click **Save**.

You must save the changes for the reorder to be displayed.

Sending a support bundle through web UI and maintenance console

You should send a support bundle when the issue that you have requires more detailed diagnosis and troubleshooting than an AutoSupport message provides. You can send a support bundle to technical support using the Unified Manager web UI and maintenance console.

Unified Manager stores a maximum of two full support bundles and three lightweight support bundles at one time.

Related information

[Unified Manager user roles and capabilities](#)

Sending AutoSupport messages and support bundles to technical support

The AutoSupport page enables you to send predefined and on-demand AutoSupport messages to your technical support team to ensure a correct operation of your environment, and to assist you in maintaining the integrity of your environment. AutoSupport is enabled by default and it should not be disabled, for you to receive the benefits of NetAppActive IQ.

You can send diagnostic system information and detailed data about the Unified Manager server in a message as and when required, schedule a message to be sent periodically, or even generate and send support bundles to the technical support team.



A user with a storage administrator role can generate and send on-demand AutoSupport messages and support bundles to technical support. However, only an administrator or maintenance user can enable or disable periodic AutoSupport and configure the HTTP settings as described in the Setting up HTTP proxy server section. In an environment that needs to use an HTTP proxy server, the configuration should be complete before a storage administrator can send on-demand AutoSupport messages and support bundles to technical support.

Sending on-demand AutoSupport messages

You can generate and send an on-demand message to technical support, or to a specified email recipient, or to both.

Steps

1. Navigate to **General > AutoSupport**, and perform one or both of the following actions:
2. If you want to send the AutoSupport message to technical support, select the **Send to Technical Support** check box.
3. If you want to send the AutoSupport message to a specific email recipient, select the **Send to Email Recipient** check box, and enter the email address of the recipient.
4. Click **Save**.
5. Click **Generate and Send AutoSupport**.

Enabling periodic AutoSupport

You can send specific, predefined messages to technical support for issue diagnosis and resolution periodically. This functionality is enabled by default. If disabled, an administrator or maintenance user can enable the settings.

Steps

1. Navigate to **General > AutoSupport**.
2. In the Periodic AutoSupport section, select the **Enable Sending AutoSupport Data Periodically to Active IQ** check box.
3. If required, define the name, port, and authentication information for the HTTP proxy server as described in Setting up HTTP proxy server section.
4. Click **Save**.

Uploading on-demand support bundle

You can generate and send a support bundle to technical support based on the requirement for troubleshooting. Unified Manager stores only the two most recently generated support bundles. Older support bundles are deleted from the system.

Because some types of support data can use a large amount of cluster resources or take a long time to complete, when you select the full support bundle, you can include or exclude specific data types to reduce the support bundle size. You also have the option to create a lightweight support bundle that contains just 30 days of logs and configuration database records — it excludes performance data, acquisition recording files, and server heap dump.

Steps

1. Navigate to **General > AutoSupport**.
2. In the On-Demand Support Bundle section, click **Generate and Send Support Bundle**.
3. To send a light support bundle to technical support, in the Generate and Send Support Bundle pop-up, select the **Generate light support bundle** check box.
4. Alternately, to send a full support bundle, select the **Generate full support bundle** check box. Select the specific data types to include or exclude in the support bundle.



Even if you do not select any data type, the support bundle is still generated with other Unified Manager data.

5. Select the **Send the bundle to technical support** check box to generate and send the bundle to technical support. If you do not select this check box, the bundle is generated and stored locally in the Unified Manager server. The generated support bundle is available for later use in the /support directory on VMware systems, in /opt/netapp/data/support/ on Linux systems, and in ProgramData\NetApp\OnCommandAppData\ocum\support on Windows systems.
6. Click **Send**.

Setting up HTTP proxy server

You can designate a proxy to provide the internet access in order to send AutoSupport content to support if your environment does not provide direct access from the Unified Manager server. This section is available for only administrator and maintenance users.

• Use HTTP proxy

Check this box to identify the server being used as the HTTP proxy.

Enter the host name or IP address of the proxy server, and the port number used to connect to the server.

• Use authentication

Check this box if you need to provide authentication information to access the server being used as the HTTP proxy.

Enter the user name and the password required to authenticate with the HTTP proxy.



HTTP proxies that provide only Basic Authentication are not supported.

Accessing the maintenance console

If the Unified Manager user interface is not in operation, or if you need to perform functions that are not available in the user interface, you can access the maintenance console to manage your Unified Manager system.

What you'll need

You must have installed and configured Unified Manager.

After 15 minutes of inactivity, the maintenance console logs you out.



When installed on VMware, if you have already logged in as the maintenance user through the VMware console, you cannot simultaneously log in using Secure Shell.

Step

1. Follow these steps to access the maintenance console:

On this operating system...	Follow these steps...
VMware	<ol style="list-style-type: none"> a. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager virtual appliance. b. Log in to the maintenance console using your maintenance user name and password.
Linux	<ol style="list-style-type: none"> a. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system. b. Log in to the system with the maintenance user (umadmin) name and password. c. Enter the command <code>maintenance_console</code> and press Enter.
Windows	<ol style="list-style-type: none"> a. Log in to the Unified Manager system with administrator credentials. b. Launch PowerShell as a Windows administrator. c. Enter the command <code>maintenance_console</code> and press Enter.

The Unified Manager maintenance console menu is displayed.

Generating and uploading a support bundle

You can generate a support bundle, containing diagnostic information, so that you can send it to technical support for troubleshooting help.

Starting with Unified Manager 9.8, if your Unified Manager server is connected to the internet, you can also upload the support bundle to NetApp from the maintenance console.

What you'll need

You must have access to the maintenance console as the maintenance user.

Because some types of support data can use a large amount of cluster resources or take a long time to complete, when you select the full support bundle you can specify data types to include or exclude to reduce the support bundle size. You also have the option to create a lightweight support bundle that contains just 30 days of logs and configuration database records — it excludes performance data, acquisition recording files, and server heap dump.

Unified Manager stores only the two most recently generated support bundles. Older support bundles are

deleted from the system.

Steps

1. In the maintenance console **Main Menu**, select **Support/Diagnostics**.
2. Select **Generate Light Support Bundle** or **Generate Support Bundle** depending on the level of details you want to have in the support bundle.
3. If you choose the full support bundle, select or deselect the following data types to include or exclude in the support bundle:

- **database dump**

A dump of the MySQL Server database.

- **heap dump**

A snapshot of the state of the main Unified Manager server processes. This option is disabled by default and should be selected only when requested by customer support.

- **acquisition recordings**

A recording of all communications between Unified Manager and the monitored clusters.



If you deselect all data types, the support bundle is still generated with other Unified Manager data.

4. Type `g`, and then press Enter to generate the support bundle.

Since the generation of a support bundle is a memory intensive operation, you are prompted to verify that you are sure you want to generate the support bundle at this time.

5. Type `y`, and then press Enter to generate the support bundle.

If you do not want to generate the support bundle at this time, type `n`, and then press Enter.

6. If you included database dump files in the full support bundle, you are prompted to specify the time period for which you want performance statistics included. Including performance statistics can take a lot of time and space, so you can also dump the database without including performance statistics:

- a. Enter the starting date in the format YYYYMMDD.

For example, enter `20210101` for January 1, 2021. Enter `n` if you do not want performance statistics to be included.

- b. Enter the number of days of statistics to include, beginning from 12 a.m. on the specified starting date.

You can enter a number from 1 through 10.

If you are including performance statistics, the system displays the period of time for which performance statistics will be collected.

7. After the support bundle is created you are prompted whether you want to upload it to NetApp. Type `y`, and then press Enter.

You are prompted to enter your support case number.

8. If you have a case number already, enter the number and press Enter. Otherwise just press Enter.

The support bundle is uploaded to NetApp.

If your Unified Manager server is not connected to the internet, or if you are unable to upload the support bundle for any other reason, then you can retrieve it and send it manually. You can retrieve it using an SFTP client or by using UNIX or Linux CLI commands. On Windows installations you can use Remote Desktop (RDP) to retrieve the support bundle.

The generated support bundle resides in the /support directory on VMware systems, in /opt/netapp/data/support/ on Linux systems, and in ProgramData\NetApp\OnCommandAppData\ocum\support on Windows systems.

Related information

[Unified Manager user roles and capabilities](#)

Retrieving the support bundle using a Windows client

If you are a Windows user, you can download and install a tool to retrieve the support bundle from your Unified Manager server. You can send the support bundle to technical support for a more detailed diagnosis of an issue. Filezilla or WinSCP are examples of tools you can use.

What you'll need

You must be the maintenance user to perform this task.

You must use a tool that supports SCP or SFTP.

Steps

1. Download and install a tool to retrieve the support bundle.
2. Open the tool.
3. Connect to your Unified Manager management server over SFTP.

The tool displays the contents of the /support directory and you can view all existing support bundles.

4. Select the destination directory for the support bundle you want to copy.
5. Select the support bundle you want to copy and use the tool to copy the file from the Unified Manager server to your local system.

Retrieving the support bundle using a UNIX or Linux client

If you are a UNIX or Linux user, you can retrieve the support bundle from your vApp by using the command-line interface (CLI) on your Linux client server. You can use either SCP or SFTP to retrieve the support bundle.

What you'll need

You must be the maintenance user to perform this task.

You must have generated a support bundle using the maintenance console and have the support bundle name

available.

Steps

1. Access the CLI through Telnet or the console, using your Linux client server.
2. Access the `/support` directory.
3. Retrieve the support bundle and copy it to the local directory using the following command:

If you are using...	Then use the following command...
SCP	<code>scp <maintenance-user>@<vApp-name-or-ip>:/support/support_bundle_file_name.7z <destination-directory></code>
SFTP	<code>sftp <maintenance-user>@<vApp-name-or-ip>:/support/support_bundle_file_name.7z <destination-directory></code>

The name of the support bundle is provided to you when you generate it using the maintenance console.

4. Enter the maintenance user password.

Examples

The following example uses SCP to retrieve the support bundle:

```
`$ scp
admin@10.10.12.69:/support/support_bundle_20160216_145359.7z .`
Password: `<maintenance_user_password>`
support_bundle_20160216_145359.7z 100% 119MB 11.9MB/s 00:10
```

The following example uses SFTP to retrieve the support bundle:

```
`$ sftp
admin@10.10.12.69:/support/support_bundle_20160216_145359.7z .`
Password: `<maintenance_user_password>`
Connected to 10.228.212.69.
Fetching /support/support_bundle_20130216_145359.7z to
./support_bundle_20130216_145359.7z
/support/support_bundle_20160216_145359.7z
```

Sending a support bundle to technical support

When an issue requires more detailed diagnosis and troubleshooting information than an AutoSupport message provides, you can send a support bundle to technical support.

What you'll need

You must have access to the support bundle to send it to technical support.

You must have a case number generated through the technical support web site.

Steps

1. Log in to the NetApp Support Site.
2. Upload the file.

[How to upload a file to NetApp](#)

Tasks and information related to several workflows

Some tasks and reference texts that can help you understand and complete a workflow are common to many of the workflows in Unified Manager, including adding and reviewing notes about an event, assigning an event, acknowledging and resolving events, and details about volumes, storage virtual machines (SVMs), aggregates, and so on.

Cluster components and why they can be in contention

You can identify cluster performance issues when a cluster component goes into contention. The performance of workloads that use the component slow down and their response time (latency) for client requests increases, which triggers an event in Unified Manager.

A component that is in contention cannot perform at an optimal level. Its performance has declined, and the performance of other cluster components and workloads, called *victims*, might have increased latency. To bring a component out of contention, you must reduce its workload or increase its ability to handle more work, so that the performance can return to normal levels. Because Unified Manager collects and analyzes workload performance in five-minute intervals, it detects only when a cluster component is consistently overused. Transient spikes of overusage that last for only a short duration within the five-minute interval are not detected.

For example, a storage aggregate might be under contention because one or more workloads on it are competing for their I/O requests to be fulfilled. Other workloads on the aggregate can be impacted, causing their performance to decrease. To reduce the amount of activity on the aggregate, there are different steps you can take, such as moving one or more workloads to a less busy aggregate or node, to lessen the overall workload demand on the current aggregate. For a QoS policy group, you can adjust the throughput limit, or move workloads to a different policy group, so that the workloads are no longer being throttled.

Unified Manager monitors the following cluster components to alert you when they are in contention:

- **Network**

Represents the wait time of I/O requests by the external networking protocols on the cluster. The wait time is time spent waiting for “transfer ready” transactions to finish before the cluster can respond to an I/O request. If the network component is in contention, it means high wait time at the protocol layer is impacting the latency of one or more workloads.

- **Network Processing**

Represents the software component in the cluster involved with I/O processing between the protocol layer and the cluster. The node handling network processing might have changed since the event was detected. If the network processing component is in contention, it means high utilization at the network processing

node is impacting the latency of one or more workloads.

When using an All SAN Array cluster in an active-active configuration, the network processing latency value is displayed for both nodes so you can verify the nodes are sharing the load equally.

- **QoS Limit Max**

Represents the throughput maximum (peak) setting of the storage Quality of Service (QoS) policy group assigned to the workload. If the policy group component is in contention, it means all workloads in the policy group are being throttled by the set throughput limit, which is impacting the latency of one or more of those workloads.

- **QoS Limit Min**

Represents the latency to a workload that is being caused by QoS throughput minimum (expected) setting assigned to other workloads. If the QoS minimum set on certain workloads use the majority of the bandwidth to guarantee the promised throughput, other workloads will be throttled and see more latency.

- **Cluster Interconnect**

Represents the cables and adapters with which clustered nodes are physically connected. If the cluster interconnect component is in contention, it means high wait time for I/O requests at the cluster interconnect is impacting the latency of one or more workloads.

- **Data Processing**

Represents the software component in the cluster involved with I/O processing between the cluster and the storage aggregate that contains the workload. The node handling data processing might have changed since the event was detected. If the data processing component is in contention, it means high utilization at the data processing node is impacting the latency of one or more workloads.

- **Volume Activation**

Represents the process that tracks the usage of all active volumes. In large environments where more than 1000 volumes are active, this process tracks how many critical volumes need to access resources through the node at the same time. When the number of concurrent active volumes exceeds the recommended maximum threshold, some of the non-critical volumes will experience latency as identified here.

- **MetroCluster Resources**

Represents the MetroCluster resources, including NVRAM and interswitch links (ISLs), used to mirror data between clusters in a MetroCluster configuration. If the MetroCluster component is in contention, it means high write throughput from workloads on the local cluster or a link health issue is impacting the latency of one or more workloads on the local cluster. If the cluster is not in a MetroCluster configuration, this icon is not displayed.

- **Aggregate or SSD Aggregate Ops**

Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate), or a mix of HDDs and a cloud tier (a FabricPool aggregate). An “SSD Aggregate” consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

- **Cloud Latency**

Represents the software component in the cluster involved with I/O processing between the cluster and the cloud tier on which user data is stored. If the cloud latency component is in contention, it means that a large amount of reads from volumes that are hosted on the cloud tier are impacting the latency of one or more workloads.

- **Sync SnapMirror**

Represents the software component in the cluster involved with replicating user data from the primary volume to the secondary volume in a SnapMirror Synchronous relationship. If the sync SnapMirror component is in contention, it means that the activity from SnapMirror Synchronous operations are impacting the latency of one or more workloads.

Volume / Health details page

You can use the Volume / Health details page to view detailed information about a selected volume, such as capacity, storage efficiency, configuration, protection, annotation, and events generated. You can also view information about the related objects and related alerts for that volume.

You must have the Application Administrator or Storage Administrator role.

Command buttons

The command buttons enable you to perform the following tasks for the selected volume:

- **Switch to Performance View**

Enables you to navigate to the Volume / Performance details page.

- **Actions**

- Add Alert

Enables you to add an alert to the selected volume.

- Edit Thresholds

Enables you to modify the threshold settings for the selected volume.

- Annotate

Enables you to annotate the selected volume.

- Protect

Enables you to create either SnapMirror or SnapVault relationships for the selected volume.

- Relationship

Enables you to execute the following protection relationship operations:

- Edit

Launches the Edit Relationship dialog box which enables you to change existing SnapMirror

policies, schedules, and maximum transfer rates for an existing protection relationship.

- Abort

Aborts transfers that are in progress for a selected relationship. Optionally, it enables you to remove the restart checkpoint for transfers other than the baseline transfer. You cannot remove the checkpoint for a baseline transfer.

- Quiesce

Temporarily disables scheduled updates for a selected relationship. Transfers that are already in progress must complete before the relationship is quiesced.

- Break

Breaks the relationship between the source and destination volumes and changes the destination to a read-write volume.

- Remove

Permanently deletes the relationship between the selected source and destination. The volumes are not destroyed and the Snapshot copies on the volumes are not removed. This operation cannot be undone.

- Resume

Enables scheduled transfers for a quiesced relationship. At the next scheduled transfer interval, a restart checkpoint is used, if one exists.

- Resynchronize

Enables you to resynchronize a previously broken relationship.

- Initialize/Update

Enables you to perform a first-time baseline transfer on a new protection relationship, or to perform a manual update if the relationship is already initialized.

- Reverse Resync

Enables you to reestablish a previously broken protection relationship, reversing the function of the source and destination by making the source a copy of the original destination. The contents on the source are overwritten by the contents on the destination, and any data that is newer than the data on the common Snapshot copy is deleted.

- Restore

Enables you to restore data from one volume to another volume. For information, see [Restoring data using the Volume / Health details page](#).



The Restore button and the Relationship operation buttons are not available for volumes that are in synchronous protection relationships.

- **View Volumes**

Enables you to navigate to the Health: All Volumes view.

Capacity tab

The Capacity tab displays details about the selected volume, such as its physical capacity, logical capacity, threshold settings, quota capacity, and information about any volume move operation:

- **Capacity Physical**

Details the physical capacity of the volume:

- Snapshot Overflow

Displays the data space that is consumed by the Snapshot copies.

- Used

Displays the space used by data in the volume.

- Warning

Indicates that the space in the volume is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

- Error

Indicates that the space in the volume is full. If this threshold is breached, the Space Full event is generated.

- Unusable

Indicates that the Thin-Provisioned Volume Space At Risk event is generated and that the space in the thinly provisioned volume is at risk because of aggregate capacity issues. The unusable capacity is displayed only for thinly provisioned volumes.

- Data graph

Displays the total data capacity and the used data capacity of the volume.

If autogrow is enabled, the data graph also displays the space available in the aggregate. The data graph displays the effective storage space that can be used by data in the volume, which can be one of the following:

- Actual data capacity of the volume for the following conditions:
 - Autogrow is disabled.
 - Autogrow-enabled volume has reached the maximum size.
 - Autogrow-enabled thickly provisioned volume cannot grow further.
- Data capacity of the volume after considering the maximum volume size (for thinly provisioned volumes and for thickly provisioned volumes when the aggregate has space for the volume to reach maximum size)
- Data capacity of the volume after considering the next possible autogrow size (for thickly provisioned volumes that have an autogrow percentage threshold)

- Snapshot copies graph

This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

- **Capacity Logical**

Displays the logical space characteristics of the volume. The logical space indicates the real size of the data that is being stored on disk without applying the savings from using ONTAP storage efficiency technologies.

- Logical Space Reporting

Displays if the volume has logical space reporting configured. The value can be Enabled, Disabled, or Not applicable. "Not applicable" is displayed for volumes on older versions of ONTAP or on volumes that do not support logical space reporting.

- Used

Displays the amount of logical space that is being used by data in the volume, and the percentage of logical space used based on the total data capacity.

- Logical Space Enforcement

Displays whether logical space enforcement is configured for thinly provisioned volumes. When set to Enabled, the logical used size of the volume cannot be greater than the currently set physical volume size.

- **Autogrow**

Displays whether the volume automatically grows when it is out of space.

- **Space Guarantee**

Displays the FlexVol volume setting control when a volume removes free blocks from an aggregate. These blocks are then guaranteed to be available for writes to files in the volume. The space guarantee can be set to one of the following:

- None

No space guarantee is configured for the volume.

- File

Full size of sparsely written files (for example, LUNs) is guaranteed.

- Volume

Full size of the volume is guaranteed.

- Partial

The FlexCache volume reserves space based on its size. If the FlexCache volume's size is 100 MB or more, the minimum space guarantee is set to 100 MB by default. If the FlexCache volume's size is less

than 100 MB, the minimum space guarantee is set to the FlexCache volume's size. If the FlexCache volume's size is grown later, the minimum space guarantee is not incremented.



The space guarantee is Partial when the volume is of type Data-Cache.

- **Details (Physical)**

Displays the physical characteristics of the volume.

- **Total Capacity**

Displays the total physical capacity in the volume.

- **Data Capacity**

Displays the amount of physical space used by the volume (used capacity) and the amount of physical space that is still available (free capacity) in the volume. These values are also displayed as a percentage of the total physical capacity.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

- **Snapshot Reserve**

Displays the amount of space used by the Snapshot copies (used capacity) and amount of space available for Snapshot copies (free capacity) in the volume. These values are also displayed as a percentage of the total snapshot reserve.

When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the Snapshot copies (used capacity) and the amount of space that is available in the volume but cannot be used for making Snapshot copies (unusable capacity) because of aggregate capacity issues is displayed.

- **Volume Thresholds**

Displays the following volume capacity thresholds:

- Nearly Full Threshold

Specifies the percentage at which a volume is nearly full.

- Full Threshold

Specifies the percentage at which a volume is full.

- **Other Details**

- Autogrow Max Size

Displays the maximum size up to which the volume can automatically grow. The default value is 120% of the volume size on creation. This field is displayed only when autogrow is enabled for the volume.

- Qtree Quota Committed Capacity

Displays the space reserved in the quotas.

- Qtree Quota Overcommitted Capacity

Displays the amount of space that can be used before the system generates the Volume Qtree Quota Overcommitted event.

- Fractional Reserve

Controls the size of the overwrite reserve. By default, the fractional reserve is set to 100, indicating that 100 percent of the required reserved space is reserved so that the objects are fully protected for overwrites. If the fractional reserve is less than 100 percent, the reserved space for all the space-reserved files in that volume is reduced to the fractional reserve percentage.

- Snapshot Daily Growth Rate

Displays the change (in percentage, or in KB, MB, GB, and so on) that occurs every 24 hours in the Snapshot copies in the selected volume.

- Snapshot Days to Full

Displays the estimated number of days remaining before the space reserved for the Snapshot copies in the volume reaches the specified threshold.

The Snapshot Days to Full field displays a Not Applicable value when the growth rate of the Snapshot copies in the volume is zero or negative, or when there is insufficient data to calculate the growth rate.

- Snapshot Autodelete

Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.

- Snapshot Copies

Displays information about the Snapshot copies in the volume.

The number of Snapshot copies in the volume is displayed as a link. Clicking the link opens the Snapshot Copies on a Volume dialog box, which displays details of the Snapshot copies.

The Snapshot copy count is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

- **Volume Move**

Displays the status of either the current or the last volume move operation that was performed on the volume, and other details, such as the current phase of the volume move operation which is in progress, source aggregate, destination aggregate, start time, end time, and estimated end time.

Also displays the number of volume move operations that are performed on the selected volume. You can view more information about the volume move operations by clicking the **Volume Move History** link.

Configuration tab

The Configuration tab displays details about the selected volume, such as the export policy, RAID type, capacity and storage efficiency related features of the volume:

• Overview

- Full Name

Displays the full name of the volume.

- Aggregates

Displays the name of the aggregate on which the volume resides, or the number of aggregates on which the FlexGroup volume resides.

- Tiering Policy

Displays the tiering policy set for the volume; if the volume is deployed on a FabricPool-enabled aggregate. The policy can be None, Snapshot Only, Backup, Auto, or All.

- Storage VM

Displays the name of the SVM that contains the volume.

- Junction Path

Displays the status of the path, which can be active or inactive. The path in the SVM to which the volume is mounted is also displayed. You can click the **History** link to view the most recent five changes to the junction path.

- Export Policy

Displays the name of the export policy that is created for the volume. You can click the link to view details about the export policies, authentication protocols, and access enabled on the volumes that belong to the SVM.

- Style

Displays the volume style. The volume style can be FlexVol or FlexGroup.

- Type

Displays the type of the selected volume. The volume type can be Read-write, Load-sharing, Data-Protection, Data-cache, or Temporary.

- RAID Type

Displays the RAID type of the selected volume. The RAID type can be RAID0, RAID4, RAID-DP, or RAID-TEC.



Multiple RAID types may display for FlexGroup volumes because the constituent volumes for FlexGroups can be on aggregates of different types.

- SnapLock Type

Displays the SnapLock Type of the aggregate that contains the volume.

- SnapLock Expiry

Displays the expiry date of SnapLock volume.

- **Capacity**

- Thin Provisioning

Displays whether thin provisioning is configured for the volume.

- Autogrow

Displays whether the flexible volume grows automatically within an aggregate.

- Snapshot Autodelete

Specifies whether Snapshot copies are automatically deleted to free space when a write to a volume fails because of lack of space in the aggregate.

- Quotas

Specifies whether the quotas are enabled for the volume.

- **Efficiency**

- Compression

Specifies whether compression is enabled or disabled.

- Deduplication

Specifies whether deduplication is enabled or disabled.

- Deduplication Mode

Specifies whether the deduplication operation enabled on a volume is a manual, scheduled, or policy-based operation. If the mode is set to Scheduled, the operation schedule is displayed, and if the mode is set to a policy, the policy name is displayed.

- Deduplication Type

Specifies the type of deduplication operation running on the volume. If the volume is in a SnapVault relationship, the type displayed is SnapVault. For any other volume, the type is displayed as Regular.

- Storage Efficiency Policy

Specifies the name of the storage efficiency policy that has been assigned through Unified Manager to this volume. This policy can control the compression and deduplication settings.

- **Protection**

- Snapshot Copies

Specifies whether automatic Snapshot copies are enabled or disabled.

Protection tab

The Protection tab displays protection details about the selected volume, such as lag information, relationship type, and topology of the relationship.

- **Summary**

Displays protection relationships (SnapMirror, SnapVault, or Storage VM DR) properties for a selected volume. For any other relationship type, only the Relationship Type property is displayed. If a primary volume is selected, only the Managed and Local Snapshot copy Policy are displayed. Properties displayed for SnapMirror and SnapVault relationships include the following:

- Source Volume

Displays the name of the selected volume's source if the selected volume is a destination.

- Lag Status

Displays the update or transfer lag status for a protection relationship. The status can be Error, Warning, or Critical.

The lag status is not applicable for synchronous relationships.

- Lag Duration

Displays the time by which the data on the mirror lags behind the source.

- Last Successful Update

Displays the date and time of the most recent successful protection update.

The last successful update is not applicable for synchronous relationships.

- Storage Service Member

Displays either Yes or No to indicate whether or not the volume belongs to and is managed by a storage service.

- Version Flexible Replication

Displays either Yes, Yes with backup option, or None. Yes indicates that SnapMirror replication is possible even if source and destination volumes are running different versions of ONTAP software. Yes with backup option indicates the implementation of SnapMirror protection with the ability to retain multiple versions of backup copies on the destination. None indicates that Version Flexible Replication is not enabled.

- Relationship Capability

Indicates the ONTAP capabilities available to the protection relationship.

- Protection Service

Displays the name of the protection service if the relationship is managed by a protection partner application.

- Relationship Type

Displays any relationship type, including Asynchronous Mirror, Asynchronous Vault, Asynchronous MirrorVault, StrictSync, and Sync.

- Relationship State

Displays the state of the SnapMirror or SnapVault relationship. The state can be Uninitialized,

SnapMirrored, or Broken-Off. If a source volume is selected, the relationship state is not applicable and is not displayed.

- Transfer Status

Displays the transfer status for the protection relationship. The transfer status can be one of the following:

- Aborting

SnapMirror transfers are enabled; however, a transfer abort operation that might include removal of the checkpoint is in progress.

- Checking

The destination volume is undergoing a diagnostic check and no transfer is in progress.

- Finalizing

SnapMirror transfers are enabled. The volume is currently in the post-transfer phase for incremental SnapVault transfers.

- Idle

Transfers are enabled and no transfer is in progress.

- In-Sync

The data in the two volumes in the synchronous relationship are synchronized.

- Out-of-Sync

The data in the destination volume is not synchronized with the source volume.

- Preparing

SnapMirror transfers are enabled. The volume is currently in the pre-transfer phase for incremental SnapVault transfers.

- Queued

SnapMirror transfers are enabled. No transfers are in progress.

- Quiesced

SnapMirror transfers are disabled. No transfer is in progress.

- Quiescing

A SnapMirror transfer is in progress. Additional transfers are disabled.

- Transferring

SnapMirror transfers are enabled and a transfer is in progress.

- Transitioning

The asynchronous transfer of data from the source to the destination volume is complete, and the transition to synchronous operation has started.

- Waiting

A SnapMirror transfer has been initiated, but some associated tasks are waiting to be queued.

- Max Transfer Rate

Displays the maximum transfer rate for the relationship. The maximum transfer rate can be a numerical value in either kilobytes per second (Kbps), Megabytes per second (Mbps), Gigabytes per second (Gbps), or Terabytes per second (Tbps). If No Limit is displayed, the baseline transfer between relationships is unlimited.

- SnapMirror Policy

Displays the protection policy for the volume. DPDefault indicates the default Asynchronous Mirror protection policy, XDPDefault indicates the default Asynchronous Vault policy, and DPSyncDefault indicates the default Asynchronous MirrorVault policy. StrictSync indicates the default Synchronous Strict protection policy, and Sync indicates the default Synchronous policy. You can click the policy name to view details associated with that policy, including the following information:

- Transfer priority
- Ignore access time setting
- Tries limit
- Comments
- SnapMirror labels
- Retention settings
- Actual Snapshot copies
- Preserve Snapshot copies
- Retention warning threshold
- Snapshot copies with no retention settings In a cascading SnapVault relationship where the source is a data protection (DP) volume, only the rule “sm_created” applies.

- Update Schedule

Displays the SnapMirror schedule assigned to the relationship. Positioning your cursor over the information icon displays the schedule details.

- Local Snapshot Policy

Displays the Snapshot copy policy for the volume. The policy is Default, None, or any name given to a custom policy.

- Protected By

Displays the type of protection used for the selected volume. For example, if a volume is protected by Consistency Group and SnapMirror volume relationships, this field displays both SnapMirror and Consistency Group. This field also provides a link that redirects you to the Relationships page to view the unified relationship status. The link is only applicable to constituent relationships.

- Consistency Group

For volumes protected by SnapMirror Business Continuity (SM-BC) relationships, this column displays the Consistency Group of the volume.

- **Views**

Displays the protection topology of the selected volume. The topology includes graphical representations of all volumes that are related to the selected volume. The selected volume is indicated by a dark gray border, and lines between volumes in the topology indicate the protection relationship type. The direction of the relationships in the topology are displayed from left to right, with the source of each relationship on the left and the destination on the right.

Double bold lines specify an Asynchronous Mirror relationship, a single bold line specifies an Asynchronous Vault relationship, double single lines specify an Asynchronous MirrorVault relationship, and a bold line and non-bold line specifies a Synchronous relationship. The table below indicates if the Synchronous relationship is StrictSync or Sync.

Right-clicking a volume displays a menu from which you can choose either to protect the volume or restore data to it. Right-clicking a relationship displays a menu from which you can choose to either edit, abort, quiesce, break, remove, or resume a relationship.

The menus will not display in the following instances:

- If RBAC settings do not allow this action, for example, if you have only operator privileges
- If the volume is in a synchronous protection relationship
- When the volume ID is unknown, for example, when you have an intercluster relationship and the destination cluster has not yet been discovered. Clicking another volume in the topology selects and displays information for that volume. A question mark (?) in the upper-left corner of a volume indicates that either the volume is missing or that it has not yet been discovered. It might also indicate that the capacity information is missing. Positioning your cursor over the question mark displays additional information, including suggestions for remedial action.

The topology displays information about volume capacity, lag, Snapshot copies, and last successful data transfer if it conforms to one of several common topology templates. If a topology does not conform to one of those templates, information about volume lag and last successful data transfer is displayed in a relationship table under the topology. In that case, the highlighted row in the table indicates the selected volume, and, in the topology view, bold lines with a blue dot indicate the relationship between the selected volume and its source volume.

Topology views include the following information:

- **Capacity**

Displays the total amount of capacity used by the volume. Positioning your cursor over a volume in the topology displays the current warning and critical threshold settings for that volume in the Current Threshold Settings dialog box. You can also edit the threshold settings by clicking the **Edit Thresholds** link in the Current Threshold Settings dialog box. Clearing the **Capacity** check box hides all capacity information for all volumes in the topology.

- **Lag**

Displays the lag duration and the lag status of the incoming protection relationships. Clearing the **Lag** check box hides all lag information for all volumes in the topology. When the **Lag** check box is dimmed,

then the lag information for the selected volume is displayed in the relationship table below the topology, as well as the lag information for all related volumes.

- **Snapshot**

Displays the number of Snapshot copies available for a volume. Clearing the **Snapshot** check box hides all Snapshot copy information for all volumes in the topology. Clicking a Snapshot copy icon () displays the Snapshot copy list for a volume. The Snapshot copy count displayed next to the icon is updated approximately every hour; however, the list of Snapshot copies is updated at the time that you click the icon. This might result in a difference between the Snapshot copy count displayed in the topology and the number of Snapshot copies listed when you click the icon.

- **Last Successful Transfer**

Displays the amount, duration, time, and date of the last successful data transfer. When the **Last Successful Transfer** check box is dimmed, then the last successful transfer information for the selected volume is displayed in the relationship table below the topology, as well as the last successful transfer information for all related volumes.

- **History**

Displays in a graph the history of incoming SnapMirror and SnapVault protection relationships for the selected volume. There are three history graphs available: incoming relationship lag duration, incoming relationship transfer duration, and incoming relationship transfer size. History information is displayed only when you select a destination volume. If you select a primary volume, the graphs are empty, and the message No data found is displayed. If the volumes are protected by Consistency Group and SnapMirror synchronous relationships, information for the relationship transfer duration and relationship transfer size is not displayed.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if large amounts of data are being transferred at the same time of the day or week, or if the lag warning or lag error threshold is consistently being breached, you can take the appropriate action. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Protection history graphs display the following information:

- **Relationship Lag Duration**

Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum lag duration reached in the duration period shown in the x axis. The horizontal orange line on the graph depicts the lag error threshold, and the horizontal yellow line depicts the lag warning threshold. Positioning your cursor over these lines displays the threshold setting. The horizontal blue line depicts the lag duration. You can view the details for specific points on the graph by positioning your cursor over an area of interest.

- **Relationship Transfer Duration**

Displays seconds, minutes, or hours on the vertical (y) axis, and displays days, months, or years on the horizontal (x) axis, depending on the selected duration period. The upper value on the y axis indicates the maximum transfer duration reached in the duration period shown in the x axis. You can view the details of specific points on the graph by positioning your cursor over the area of interest.



This chart is not available for volumes that are in synchronous protection relationships.

- **Relationship Transferred Size**

Displays bytes, kilobytes, megabytes, and so on, on the vertical (y) axis depending on the transfer size, and displays days, months, or years on the horizontal (x) axis depending on the selected time period. The upper value on the y axis indicates the maximum transfer size reached in the duration period shown in the x axis. You can view the details for specific points on the graph by positioning your cursor over an area of interest.



This chart is not available for volumes that are in synchronous protection relationships.

History area

The History area displays graphs that provide information about the capacity and space reservations of the selected volume. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

Graphs might be empty and the message No data found displayed when the data or the state of the volume remains unchanged for a period of time.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends—for example, if the volume usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

- **Volume Capacity Used**

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the Volume Used Capacity graph line is hidden.

- **Volume Capacity Used vs Total**

Displays the trend in how volume capacity is used based on the usage history, as well as the used capacity, total capacity, and details of the space savings from deduplication and compression, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

- **Volume Capacity Used (%)**

Displays the used capacity in the volume and the trend in how volume capacity is used based on the usage history, as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Volume Used Capacity legend, the

Volume Used Capacity graph line is hidden.

- **Snapshot Capacity Used (%)**

Displays the Snapshot reserve and Snapshot warning threshold as line graphs, and the capacity used by the Snapshot copies as an area graph, in percentage, on the vertical (y) axis. The Snapshot overflow is represented with different colors. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Snapshot Reserve legend, the Snapshot Reserve graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

- **Severity**

Displays the severity of the event.

- **Event**

Displays the event name.

- **Triggered Time**

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp when the event was generated is displayed.

Related Annotations pane

The Related Annotations pane enables you to view annotation details associated with the selected volume. The details include the annotation name and the annotation values that are applied to the volume. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view and navigate to the SVMs, aggregates, qtrees, LUNs, and Snapshot copies that are related to the volume:

- **Storage Virtual Machine**

Displays the capacity and the health status of the SVM that contains the selected volume.

- **Aggregate**

Displays the capacity and the health status of the aggregate that contains the selected volume. For FlexGroup volumes, the number of aggregates that comprise the FlexGroup is listed.

- **Volumes in the Aggregate**

Displays the number and capacity of all the volumes that belong to the parent aggregate of the selected volume. The health status of the volumes is also displayed, based on the highest severity level. For example, if an aggregate contains ten volumes, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical. This component does not appear for

FlexGroup volumes.

- **Qtrees**

Displays the number of qtrees that the selected volume contains and the capacity of qtrees with quota that the selected volume contains. The capacity of the qtrees with quota is displayed in relation to the volume data capacity. The health status of the qtrees is also displayed, based on the highest severity level. For example, if a volume has ten qtrees, five with Warning status and the remaining five with Critical status, then the status displayed is Critical.

- **NFS Shares**

Displays the number and status of the NFS shares associated with the volume.

- **SMB Shares**

Displays the number and status of the SMB/CIFS shares.

- **LUNs**

Displays the number and total size of all the LUNs in the selected volume. The health status of the LUNs is also displayed, based on the highest severity level.

- **User and Group Quotas**

Displays the number and status of the user and user group quotas associated with the volume and its qtrees.

- **FlexClone Volumes**

Displays the number and capacity of all the cloned volumes of the selected volume. The number and capacity are displayed only if the selected volume contains any cloned volumes.

- **Parent Volume**

Displays the name and capacity of the parent volume of a selected FlexClone volume. The parent volume is displayed only if the selected volume is a FlexClone volume.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected volume.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected volume. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Storage VM / Health details page

You can use the Storage VM / Health details page to view detailed information about the selected storage VM, such as its health, capacity, configuration, data policies, logical interfaces (LIFs), LUNs, qtrees, user, user group quotas, and protection details . You can also view information about the related objects and related alerts for the storage VM.



You can monitor only data storage VM.

Command buttons

The command buttons enable you to perform the following tasks for the selected storage VM:

- **Switch to Performance View**

Enables you to navigate to the Storage VM / Performance details page.

- **Actions**

- Add Alert

Enables you to add an alert to the selected storage VM.

- Annotate

Enables you to annotate the selected storage VM.

- **View Storage VMs**

Enables you to navigate to the Health: All Storage VMs view.

Health tab

The Health tab displays detailed information about data availability, data capacity, and protection issues of various objects such as volumes, aggregates, NAS LIFs, SAN LIFs, LUNs, protocols, services, NFS shares, and CIFS shares.

You can click the graph of an object to view the filtered list of objects. For example, you can click the volume capacity graph that displays warnings to view the list of volumes that have capacity issues with severity as warning.

- **Availability Issues**

Displays, as a graph, the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the storage VM. For example, information is displayed about the NAS LIFs and the SAN LIFs that are down and volumes that are offline.

You can also view information about the related protocols and services that are currently running, and the number and status of NFS and CIFS shares.

- **Capacity Issues**

Displays, as a graph, the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the storage VM. For example, information is displayed about aggregates that are likely to breach the set threshold values.

- **Protection Issues**

Provides a quick overview of storage VM protection-related health by displaying, as a field dialog box, the total number of relationships, including relationships that have protection issues and relationships that do not have any protection-related issues. You can also view the status of the storage VM DR relationship for the selected storage VM. The storage VM DR relationships events are displayed here and clicking on the events takes you to the event details page. When unprotected volumes exist, clicking on the link takes you to the Health: All Volumes view where you can view a filtered list of the unprotected volumes on the storage VM. The colors in the graph represent the different severity levels of the issues. Clicking a graph takes you to the Relationship: All Relationships view, where you can view a filtered list of protection relationship details. The information below the graph provides details about protection issues that can impact or have already impacted the protection of data in the storage VM. For example, information is displayed about volumes that have a Snapshot copy reserve that is almost full or about SnapMirror relationship lag issues.

Capacity tab

The Capacity tab displays detailed information about the data capacity of the selected SVM.

The following information is displayed for an Storage VM with FlexVol volume or FlexGroup volume:

- **Capacity**

The Capacity area displays details about the used and available capacity allocated from all volumes:

- Total Capacity

Displays the total capacity of the Storage VM.

- Used

Displays the space used by data in the volumes that belong to the Storage VM.

- Guaranteed Available

Displays the guaranteed available space for data that is available for volumes in the Storage VM.

- Unguaranteed

Displays the available space remaining for data that is allocated for thinly provisioned volumes in the Storage VM.

- **Volumes with Capacity Issues**

The Volumes with Capacity Issues list displays, in tabular format, details about the volumes that have capacity issues:

- Status

Indicates that the volume has a capacity-related issue of an indicated severity.

You can move the pointer over the status to view more information about the capacity-related event or events generated for the volume.

If the status of the volume is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use the **View Details** button to view more information about the event.

If the status of the volume is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.



A volume can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a volume has two events with severities of Error and Warning, only the Error severity is displayed.

- Volume

Displays the name of the volume.

- Used Data Capacity

Displays, as a graph, information about the volume capacity usage (in percentage).

- Days to Full

Displays the estimated number of days remaining before the volume reaches full capacity.

- Thin Provisioned

Displays whether space guarantee is set for the selected volume. Valid values are Yes and No.

- Aggregates

For FlexVol volumes, displays the name of the aggregate that contains the volume. For FlexGroup volumes, displays the number of aggregates that are used in the FlexGroup.

Configuration tab

The Configuration tab displays configuration details about the selected storage VM, such as its cluster, root volume, the type of volumes it contains (FlexVol volumes), policies, and protection created on the storage VM:

- **Overview**

- Cluster

Displays the name of the cluster to which the storage VM belongs.

- Allowed Volume Type

Displays the type of volumes that can be created in the storage VM. The type can be FlexVol or FlexVol/FlexGroup.

- Root Volume

Displays the name of the root volume of the storage VM.

- Allowed Protocols

Displays the type of protocols that can be configured on the storage VM. Also, indicates if a protocol is up (●), down (●), or is not configured (●).

- **Data Network Interfaces**

- NAS

Displays the number of NAS interfaces that are associated with the storage VM. Also, indicates if the interfaces are up (●) or down (●).

- SAN

Displays the number of SAN interfaces that are associated with the storage VM. Also, indicates if the interfaces are up (●) or down (●).

- FC-NVMe

Displays the number of FC-NVMe interfaces that are associated with the Storage VM. Also, indicates if the interfaces are up (●) or down (●).

- **Management Network Interfaces**

- Availability

Displays the number of management interfaces that are associated with the Storage VM. Also, indicates if the management interfaces are up (●) or down (●).

- **Policies**

- Snapshots

Displays the name of the Snapshot policy that is created on the Storage VM.

- Export Policies

Displays either the name of the export policy if a single policy is created or displays the number of export policies if multiple policies are created.

- **Protection**

- Storage VM DR

Displays whether the selected storage VM is protected, destination, or unprotected and the name of the destination on which the storage VM is protected. If the selected storage VM is destination, then the details of source storage VM are displayed. In case of fan-out, this field displays the number of total destination storage VMs on which the storage VM is protected. The count link takes you to the storage VM relationship grid filtered on source storage VM.

- Protected Volumes

Displays the number of protected volumes on the selected storage VM out of the total volumes. If you are viewing a destination storage VM, then the number link is for the destination volumes of the selected storage VM.

- Unprotected Volumes

Displays the number of unprotected volumes on the selected storage VM.

- **Services**

- Type

Displays the type of service that is configured on the storage VM. The type can be Domain Name System (DNS) or Network Information Service (NIS).

- **State**

Displays the state of the service, which can be Up (●), Down (●), or Not Configured (●).

- **Domain Name**

Displays the fully qualified domain names (FQDNs) of the DNS server for the DNS services or NIS server for the NIS services. When the NIS server is enabled, the active FQDN of the NIS server is displayed. When the NIS server is disabled, the list of all the FQDNs are displayed.

- **IP Address**

Displays the IP addresses of the DNS or NIS server. When the NIS server is enabled, the active IP address of the NIS server is displayed. When the NIS server is disabled, the list of all the IP addresses are displayed.

Network Interfaces tab

The Network Interfaces tab displays details about the data network interfaces (LIFs) that are created on the selected storage VM:

- **Network Interface**

Displays the name of the interface that is created on the selected storage VM.

- **Operational Status**

Displays the operational status of the interface, which can be Up (↑), Down (↓), or Unknown (?). The operational status of an interface is determined by the status of its physical ports.

- **Administrative Status**

Displays the administrative status of the interface, which can be Up (↑), Down (↓), or Unknown (?). The administrative status of an interface is controlled by the storage administrator to make changes to the configuration or for maintenance purposes. The administrative status can be different from the operational status. However, if the administrative status of an interface is Down, the operational status is Down by default.

- **IP Address / WWPN**

Displays the IP address for Ethernet interfaces and the World Wide Port Name (WWPN) for FC LIFs.

- **Protocols**

Displays the list of data protocols that are specified for the interface, such as CIFS, NFS, iSCSI, FC/FCoE, FC-NVMe, and FlexCache.

- **Role**

Displays the interface role. The roles can be Data or Management.

- **Home Port**

Displays the physical port to which the interface was originally associated.

- **Current Port**

Displays the physical port to which the interface is currently associated. If the interface is migrated, the current port might be different from the home port.

- **Port Set**

Displays the port set to which the interface is mapped.

- **Failover Policy**

Displays the failover policy that is configured for the interface. For NFS, CIFS, and FlexCache interfaces, the default failover policy is Next Available. Failover policy is not applicable for FC and iSCSI interfaces.

- **Routing Groups**

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

- **Failover Group**

Displays the name of the failover group.

Qtrees tab

The Qtrees tab displays details about qtrees and their quotas. You can click the **Edit Thresholds** button if you want to edit the health threshold settings for qtree capacity for one or more qtrees.

Use the **Export** button to create a comma-separated values (.csv) file containing the details of all the monitored qtrees. When exporting to a CSV file you can choose to create a qtrees report for the current storage VM, for all storage VMs in the current cluster, or for all storage VMs for all clusters in your data center. Some additional qtrees fields appear in the exported CSV file.

- **Status**

Displays the current status of the qtree. The status can be Critical () , Error () , Warning () , or Normal () .

You can move the pointer over the status icon to view more information about the event or events generated for the qtree.

If the status of the qtree is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the qtree is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.



A qtree can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a qtree has two events with severities of Error and Warning, only the Error severity is displayed.

- **Qtree**

Displays the name of the qtree.

- **Cluster**

Displays the name of the cluster containing the qtree. Appears only in the exported CSV file.

- **Storage Virtual Machine**

Displays the storage virtual machine (SVM) name containing the qtree. Appears only in the exported CSV file.

- **Volume**

Displays the name of the volume that contains the qtree.

You can move the pointer over the volume name to view more information about the volume.

- **Quota Set**

Indicates whether a quota is enabled or disabled on the qtree.

- **Quota Type**

Specifies if the quota is for a user, user group, or a qtree. Appears only in the exported CSV file.

- **User or Group**

Displays the name of the user or user group. There will be multiple rows for each user and user group. When the quota type is qtree or if the quota is not set, then the column is empty. Appears only in the exported CSV file.

- **Disk Used %**

Displays the percentage of disk space used. If a disk hard limit is set, this value is based on the disk hard limit. If the quota is set without a disk hard limit, the value is based on the volume data space. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then “Not applicable” is displayed in the grid page and the field is blank in the CSV export data.

- **Disk Hard Limit**

Displays the maximum amount of disk space allocated for the qtree. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

- **Disk Soft Limit**

Displays the amount of disk space allocated for the qtree before a warning event is generated. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk soft limit, if the quota

is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

- **Disk Threshold**

Displays the threshold value set on the disk space. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a disk threshold limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

- **Files Used %**

Displays the percentage of files used in the qtree. If the file hard limit is set, this value is based on the file hard limit. No value is displayed if the quota is set without a file hard limit. If the quota is not set or if quotas are off on the volume to which the qtree belongs, then “Not applicable” is displayed in the grid page and the field is blank in the CSV export data.

- **File Hard Limit**

Displays the hard limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file hard limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs.

- **File Soft Limit**

Displays the soft limit for the number of files permitted on the qtrees. The value is displayed as “Unlimited” for the following conditions: if the quota is set without a file soft limit, if the quota is not set, or if quotas are off on the volume to which the qtree belongs. By default, this column is hidden.

User and Group Quotas tab

Displays details about the user and user group quotas for the selected storage VM. You can view information such as the status of the quota, name of the user or user group, soft and hard limits set on the disks and files, amount of disk space and number of files used, and the disk threshold value. You can also change the email address associated with a user or user group.

- **Edit Email Address command button**

Opens the Edit Email Address dialog box, which displays the current email address of the selected user or user group. You can modify the email address. If the **Edit Email Address** field is blank, the default rule is used to generate an email address for the selected user or user group.

If more than one user has the same quota, the names of the users are displayed as comma-separated values. Also, the default rule is not used to generate the email address; therefore, you must provide the required email address for notifications to be sent.

- **Configure Email Rules command button**

Enables you to create or modify rules to generate an email address for the user or user group quotas that are configured on the storage VM. A notification is sent to the specified email address when there is a quota breach.

- **Status**

Displays the current status of the quota. The status can be Critical () , Warning () , or Normal () .

You can move the pointer over the status icon to view more information about the event or events

generated for the quota.

If the status of the quota is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can use **View Details** to view more information about the event.

If the status of the quota is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events were triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also use **View All Events** to view the list of generated events.



A quota can have multiple events of the same severity or different severities. However, only the highest severity is displayed. For example, if a quota has two events with severities of Error and Warning, only the Error severity is displayed.

- **User or Group**

Displays the name of the user or user group. If more than one user has the same quota, the names of the users are displayed as comma-separated values.

The value is displayed as “Unknown” when ONTAP does not provide a valid user name because of SecD errors.

- **Type**

Specifies if the quota is for a user or a user group.

- **Volume or Qtree**

Displays the name of the volume or qtree on which the user or user group quota is specified.

You can move the pointer over the name of the volume or qtree to view more information about the volume or qtree.

- **Disk Used %**

Displays the percentage of disk space used. The value is displayed as “Not applicable” if the quota is set without a disk hard limit.

- **Disk Hard Limit**

Displays the maximum amount of disk space allocated for the quota. Unified Manager generates a critical event when this limit is reached and no further disk writes are allowed. The value is displayed as “Unlimited” if the quota is set without a disk hard limit.

- **Disk Soft Limit**

Displays the amount of disk space allocated for the quota before a warning event is generated. The value is displayed as “Unlimited” if the quota is set without a disk soft limit. By default, this column is hidden.

- **Disk Threshold**

Displays the threshold value set on the disk space. The value is displayed as “Unlimited” if the quota is set without a disk threshold limit. By default, this column is hidden.

- **Files Used %**

Displays the percentage of files used in the qtree. The value is displayed as “Not applicable” if the quota is set without a file hard limit.

- **File Hard Limit**

Displays the hard limit for the number of files permitted on the quota. The value is displayed as “Unlimited” if the quota is set without a file hard limit.

- **File Soft Limit**

Displays the soft limit for the number of files permitted on the quota. The value is displayed as “Unlimited” if the quota is set without a file soft limit. By default, this column is hidden.

- **Email Address**

Displays the email address of the user or user group to which notifications are sent when there is a breach in the quotas.

NFS Shares tab

The NFS Shares tab displays information about NFS shares such as its status, the path associated with the volume (FlexGroup volumes or FlexVol volumes), access levels of clients to the NFS shares, and the export policy defined for the volumes that are exported. NFS shares will not be displayed in the following conditions: if the volume is not mounted or if the protocols associated with the export policy for the volume do not contain NFS shares.

- **Status**

Displays the current status of the NFS shares. The status can be Error () or Normal ()

- **Junction Path**

Displays the path to which the volume is mounted. If an explicit NFS exports policy is applied to a qtree, the column displays the path of the volume through which the qtree can be accessed.

- **Junction Path Active**

Displays whether the path to access the mounted volume is active or inactive.

- **Volume or Qtree**

Displays the name of the volume or qtree to which the NFS export policy is applied. If an NFS export policy is applied to a qtree in the volume, the column displays both the names of the volume and the qtree.

You can click the link to view details about the object in the respective details page. If the object is a qtree, links are displayed for both the qtree and the volume.

- **Volume State**

Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, or Mixed.

- Offline

Read or write access to the volume is not allowed.

- Online

Read and write access to the volume is allowed.

- Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

- Mixed

The constituents of a FlexGroup volume are not all in the same state.

- **Security Style**

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

- UNIX (NFS clients)

Files and directories in the volume have UNIX permissions.

- Unified

Files and directories in the volume have a unified security style.

- NTFS (CIFS clients)

Files and directories in the volume have Windows NTFS permissions.

- Mixed

Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

- **UNIX Permission**

Displays the UNIX permission bits in an octal string format, which is set for the volumes that are exported. It is similar to the UNIX style permission bits.

- **Export Policy**

Displays the rules that define the access permission for volumes that are exported. You can click the link to view details about the rules associated with the export policy such as the authentication protocols and the access permission.

SMB Shares tab

Displays information about the SMB shares on the selected storage VM. You can view information such as the status of the SMB share, share name, path associated with the storage VM, the status of the junction path of the share, containing object, state of the containing volume, security data of the share, and export policies defined for the share. You can also determine whether an equivalent NFS path for the SMB share exists.



Shares in folders are not displayed in the SMB Shares tab.

- **View User Mapping command button**

Launches the User Mapping dialog box.

You can view the details of user mapping for the storage VM.

- **Show ACL command button**

Launches the Access Control dialog box for the share.

You can view user and permission details for the selected share.

- **Status**

Displays the current status of the share. The status can be Normal (✓) or Error (⚠).

- **Share Name**

Displays the name of the SMB share.

- **Path**

Displays the junction path on which the share is created.

- **Junction Path Active**

Displays whether the path to access the share is active or inactive.

- **Containing Object**

Displays the name of the containing object to which the share belongs. The containing object can be a volume or a qtree.

By clicking the link, you can view details about the containing object in the respective Details page. If the containing object is a qtree, links are displayed for both qtree and volume.

- **Volume State**

Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, or Mixed.

- Offline

Read or write access to the volume is not allowed.

- Online

Read and write access to the volume is allowed.

- Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

- Mixed

The constituents of a FlexGroup volume are not all in the same state.

- **Security**

Displays the access permission for the volumes that are exported. The security style can be UNIX, Unified, NTFS, or Mixed.

- UNIX (NFS clients)

Files and directories in the volume have UNIX permissions.

- Unified

Files and directories in the volume have a unified security style.

- NTFS (CIFS clients)

Files and directories in the volume have Windows NTFS permissions.

- Mixed

Files and directories in the volume can have either UNIX permissions or Windows NTFS permissions.

- **Export Policy**

Displays the name of the export policy applicable to the share. If an export policy is not specified for the storage VM, the value is displayed as Not Enabled.

You can click the link to view details about the rules associated with the export policy, such as access protocols and permissions. The link is disabled if the export policy is disabled for the selected storage VM.

- **NFS Equivalent**

Specifies whether there is an NFS equivalent for the share.

SAN tab

Displays details about LUNs, initiator groups, and initiators for the selected storage VM. By default, the LUNs view is displayed. You can view details about the initiator groups in the Initiator Groups tab and details about initiators in the Initiators tab.

- **LUNs tab**

Displays details about the LUNs that belong to the selected storage VM. You can view information such as the LUN name, LUN state (online or offline), the name of the file system (volume or qtree) that contains the LUN, the type of host operating system, the total data capacity and serial number of the LUN. The LUN Performance column provides a link to the LUN/Performance details page.

You can also view information whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group. If it is mapped to an initiator, you can view the initiator groups and initiators that are mapped to the selected LUN.

- **Initiator Groups tab**

Displays details about initiator groups. You can view details such as the name of the initiator group, the access state, the type of host operating system that is used by all the initiators in the group, and the supported protocol. When you click the link in the access state column, you can view the current access

state of the initiator group.

- **Normal**

The initiator group is connected to multiple access paths.

- **Single Path**

The initiator group is connected to a single access path.

- **No Paths**

There is no access path connected to the initiator group.

You can view whether initiator groups are mapped to all the interfaces or specific interfaces through a port set. When you click the count link in the Mapped interfaces column, either all interfaces are displayed or specific interfaces for a port set are displayed. Interfaces that are mapped through the target portal are not displayed. The total number of initiators and LUNs that are mapped to an initiator group is displayed.

You can also view the LUNs and initiators that are mapped to the selected initiator group.

- **Initiators tab**

Displays the name and type of the initiator and the total number of initiator groups mapped to this initiator for the selected storage VM.

```
initiator groups that are mapped to the selected initiator group.
```

Related Annotations pane

The Related Annotations pane enables you to view the annotation details associated with the selected storage VM. Details include the annotation name and the annotation values that are applied to the storage VM. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

The Related Devices pane enables you to view the cluster, aggregates, and volumes that are related to the storage VM:

- **Cluster**

Displays the health status of the cluster to which the storage VM belongs.

- **Aggregates**

Displays the number of aggregates that belong to the selected storage VM. The health status of the aggregates is also displayed, based on the highest severity level. For example, if a storage VM contains ten aggregates, five of which display the Warning status and the remaining five display the Critical status, then the status displayed is Critical.

- **Assigned Aggregates**

Displays the number of aggregates that are assigned to an storage VM. The health status of the

aggregates is also displayed, based on the highest severity level.

- **Volumes**

Displays the number and capacity of the volumes that belong to the selected storage VM. The health status of the volumes is also displayed, based on the highest severity level. When there are FlexGroup volumes in the storage VM, the count also includes FlexGroups; it does not include FlexGroup constituents.

Related Groups pane

The Related Groups pane enables you to view the list of groups associated with the selected storage VM.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected storage VM. You can also add an alert by clicking the **Add Alert** link or edit an existing alert by clicking the alert name.

Cluster / Health details page

The Cluster / Health details page provides detailed information about a selected cluster, such as health, capacity, and configuration details. You can also view information about the network interfaces (LIFs), nodes, disks, related devices, and related alerts for the cluster.

The status next to the cluster name, for example (Good), represents the communication status; whether Unified Manager can communicate with the cluster. It does not represent the failover status or overall status of the cluster.

Command buttons

The command buttons enable you to perform the following tasks for the selected cluster:

- **Switch to Performance View**

Enables you to navigate to the Cluster / Performance details page.

- **Actions**

- Add Alert: Opens the Add Alert dialog box, which enables you to add an alert to the selected cluster.
- Rediscover: Initiates a manual refresh of the cluster, which enables Unified Manager to discover recent changes to the cluster.

If Unified Manager is paired with OnCommand Workflow Automation, the rediscovery operation also reacquires cached data from WFA, if any.

After the rediscovery operation is initiated, a link to the associated job details is displayed to enable tracking of the job status.

- Annotate: Enables you to annotate the selected cluster.

- **View Clusters**

Enables you to navigate to the Health: All Clusters view.

Health tab

Displays detailed information about the data availability and data capacity issues of various cluster objects such as nodes, SVMs, and aggregates. Availability issues are related to the data-serving capability of the cluster objects. Capacity issues are related to the data-storing capability of the cluster objects.

You can click the graph of an object to view a filtered list of the objects. For example, you can click the SVM capacity graph that displays warnings to view a filtered list of SVMs. This list contains SVMs that have volumes or qtrees that have capacity issues with a severity level of Warning. You can also click the SVMs availability graph that displays warnings to view the list of SVMs that have availability issues with a severity level of Warning.

Availability Issues

Graphically displays the total number of objects, including objects that have availability issues and objects that do not have any availability-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about availability issues that can impact or have already impacted the availability of data in the cluster. For example, information is displayed about disk shelves that are down and aggregates that are offline.



The data displayed for the SFO bar graph is based on the HA state of the nodes. The data displayed for all other bar graphs is calculated based on the events generated.

Capacity Issues

Graphically displays the total number of objects, including objects that have capacity issues and objects that do not have any capacity-related issues. The colors in the graph represent the different severity levels of the issues. The information below the graph provides details about capacity issues that can impact or have already impacted the capacity of data in the cluster. For example, information is displayed about aggregates that are likely to breach the set threshold values.

Capacity tab

Displays detailed information about the capacity of the selected cluster.

Capacity

Displays the data capacity graph about the used capacity and available capacity from all allocated aggregates:

- Logical Space Used

The real size of the data that is being stored on all aggregates on this cluster without applying the savings from using ONTAP storage efficiency technologies. This does not include Snapshot copies.

- Data Reduction

Displays the ratio without Snapshot copies and with two significant digits, for example, 1.8 to 1. This ratio is based on the configured ONTAP storage efficiency settings.

- Used

The physical capacity that is used by data on all aggregates. This does not include the capacity that is used for parity, right-sizing, and reservation.

- Available

Displays the capacity available for data.

- Spares

Displays the storable capacity available for storage in all the spare disks.

- Provisioned

Displays the capacity that is provisioned for all the underlying volumes.

Details

Displays detailed information about the used and available capacity. The calculation excludes the root aggregate data.

- Total Capacity

Displays the total capacity of the cluster. This does not include the capacity that is assigned for parity.

- Used

Displays the capacity that is used by data. This does not include the capacity that is used for parity, right-sizing, and reservation.

- Available

Displays the capacity available for data.

- Provisioned

Displays the capacity that is provisioned for all the underlying volumes.

- Spares

Displays the storable capacity available for storage in all the spare disks.

Cloud Tier

Displays the total cloud tier capacity used, and the capacity used for each connected cloud tier for FabricPool-enabled aggregates on the cluster. A FabricPool can be either licensed or unlicensed.

Physical Capacity Breakout by Disk Type

The Physical Capacity Breakout by Disk Type area displays detailed information about the disk capacity of the various types of disks in the cluster. By clicking the disk type, you can view more information about the disk type from the Disks tab.

- Total Usable Capacity

Displays the available capacity and spare capacity of the data disks.

- HDD

Graphically displays the used capacity and available capacity of all the HDD data disks in the cluster. The dotted line represents the spare capacity of the data disks in the HDD.

- Flash

- SSD Data

Graphically displays the used capacity and available capacity of the SSD data disks in the cluster.

- SSD Cache

Graphically displays the storable capacity of the SSD cache disks in the cluster.

- SSD Spare

Graphically displays the spare capacity of the SSD, data, and cache disks in the cluster.

- Unassigned Disks

Displays the number of unassigned disks in the cluster.

Aggregates with Capacity Issues list

Displays in tabular format details about the used capacity and available capacity of the aggregates that have capacity risk issues.

- Status

Indicates that the aggregate has a capacity-related issue of a certain severity.

You can move the pointer over the status to view more information about the event or events generated for the aggregate.

If the status of the aggregate is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. You can click the **View Details** button to view more information about the event.

If the status of the aggregate is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.



An aggregate can have multiple capacity-related events of the same severity or different severities. However, only the highest severity is displayed. For example, if an aggregate has two events with severity levels of Error and Critical, only the Critical severity is displayed.

- Aggregate

Displays the name of the aggregate.

- Used Data Capacity

Graphically displays information about the aggregate capacity usage (in percentage).

- Days to Full

Displays the estimated number of days remaining before the aggregate reaches full capacity.

Configuration tab

Displays details about the selected cluster, such as IP address, contact, and location:

Cluster Overview

- Management Interface

Displays the cluster-management LIF that Unified Manager uses to connect to the cluster. The operational status of the interface is also displayed.

- Host Name or IP Address

Displays the FQDN, short name, or the IP address of the cluster-management LIF that Unified Manager uses to connect to the cluster.

- FQDN

Displays the fully qualified domain name (FQDN) of the cluster.

- OS Version

Displays the ONTAP version that the cluster is running. If the nodes in the cluster are running different versions of ONTAP, then the earliest ONTAP version is displayed.

- Contact

Displays details about the administrator whom you should contact in case of issues with the cluster.

- Location

Displays the location of the cluster.

- Personality

Identifies if this is an All SAN Array configured cluster.

Remote Cluster Overview

Provides details about the remote cluster in a MetroCluster configuration. This information is displayed only for MetroCluster configurations.

- Cluster

Displays the name of the remote cluster. You can click the cluster name to navigate to the details page of the cluster.

- Host name or IP Address

Displays the FQDN, short name, or IP address of the remote cluster.

- Location

Displays the location of the remote cluster.

MetroCluster Overview

Provides details about the local cluster in a MetroCluster over FC or MetroCluster over IP configurations. This information is displayed only for MetroCluster over FC or IP configurations.

- Type

Displays whether the MetroCluster type is two-node or four-node. For MetroCluster over IP, only four-node is supported.

- Configuration

Displays the MetroCluster configuration over FC and IP, which can have the following values:

For FC

- Stretch Configuration with SAS cables
- Stretch Configuration with FC-SAS bridge
- Fabric Configuration with FC switches



For a four-node MetroCluster, only Fabric Configuration with FC switches is supported.

For IP

- IP configuration with Ethernet switches (L2 or L3, depending on how the cluster is configured)
 - Automated Unplanned Switch Over (AUSO)

Displays whether automated unplanned switchover is enabled for the local cluster. By default, AUSO is enabled for all clusters in a two-node MetroCluster configuration in Unified Manager. You can use the command-line interface to change the AUSO setting. This is supported for only MetroCluster over FC.

- Switch-Over mode

Displays the switch-over mode for the MetroCluster over IP configuration. The available values are: `Active`, `Negotiated Switchover`, and `Automatic Unplanned Switchover`.

Nodes

- Availability

Displays the number of nodes that are up (●) or down (●) in the cluster.

- OS Versions

Displays the ONTAP versions that the nodes are running as well as the number of nodes running a particular version of ONTAP. For example, `9.6 (2)`, `9.3 (1)` specifies that two nodes are running ONTAP 9.6, and one node is running ONTAP 9.3.

Storage Virtual Machines

- Availability

Displays the number of SVMs that are up (●) or down (●) in the cluster.

Network Interfaces

- Availability

Displays the number of non-data LIFs that are up (●) or down (●) in the cluster.

- Cluster-Management Interfaces

Displays the number of cluster-management LIFs.

- Node-Management Interfaces

Displays the number of node-management LIFs.

- Cluster Interfaces

Displays the number of cluster LIFs.

- Intercluster Interfaces

Displays the number of intercluster LIFs.

Protocols

- Data Protocols

Displays the list of licensed data protocols that are enabled for the cluster. The data protocols include iSCSI, CIFS, NFS, NVMe, and FC/FCoE.

Protection

- Mediators

Displays whether the cluster supports mediators and the connectivity status of the mediator. It indicates whether the mediator is configured, and if configured, it displays the status of the mediators.

- Not Applicable

Displays when the cluster doesn't support mediators.

- Not Configured

Displays when the cluster supports mediators, but the mediator is not configured.

- IP Address

Displays when the cluster supports mediators and the mediator is configured. The mediator status is indicated by color. The color green indicates the mediator status is reachable. The color red indicates

the mediator status is unreachable.

Cloud Tiers

Lists the names of the cloud tiers to which this cluster is connected. It also lists the type (Amazon S3, Microsoft Azure Cloud, IBM Cloud Object Storage, Google Cloud Storage, Alibaba Cloud Object Storage, or StorageGRID), and the states of the cloud tiers (Available or Unavailable).

MetroCluster Connectivity tab

Displays the issues and connectivity status of the cluster components in the MetroCluster over FC configuration. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.



The MetroCluster Connectivity tab is displayed only for clusters that are in a MetroCluster over FC configuration.

You can navigate to the details page of a remote cluster by clicking the name of the remote cluster. You can also view the details of the components by clicking the count link of a component. For example, clicking the count link of the node in the cluster displays the node tab in the details page of the cluster. Clicking the count link of the disks in the remote cluster displays the disk tab in the details page of the remote cluster.



When managing an eight-node MetroCluster configuration, clicking the count link of the Disk Shelves component displays only the local shelves of the default HA pair. Also, there is no way to display the local shelves on the other HA pair.

You can move the pointer over the components to view the details and the connectivity status of the clusters in case of any issue and to view more information about the event or events generated for the issue.

If the status of the connectivity issue between components is determined by a single event, you can view information such as the event name, time and date when the event was triggered, the name of the administrator to whom the event is assigned, and the cause of the event. The View Details button provides more information about the event.

If status of the connectivity issue between components is determined by multiple events of the same severity, the top three events are displayed with information such as the event name, time and date when the events are triggered, and the name of the administrator to whom the event is assigned. You can view more details about each of these events by clicking the event name. You can also click the **View All Events** link to view the list of generated events.

MetroCluster Replication tab

Displays the status of the data that is being replicated in a MetroCluster over FC configuration. You can use the MetroCluster Replication tab to ensure data protection by synchronously mirroring the data with the already peered clusters. A cluster is displayed in a red box when the disaster recovery partner of the cluster has issues.



The MetroCluster Replication tab is displayed only for clusters that are in a MetroCluster over FC configuration.

In a MetroCluster environment, you can use this tab to verify the logical connections and peering of the local cluster with the remote cluster. You can view the objective representation of the cluster components with their logical connections. This helps to identify the issues that might occur during mirroring of metadata and data.

In the MetroCluster Replication tab, local cluster provides the detailed graphical representation of the selected cluster and MetroCluster partner refers to the remote cluster.

Network Interfaces tab

Displays details about all the non-data LIFs that are created on the selected cluster.

Network Interface

Displays the name of the LIF that is created on the selected cluster.

Operational Status

Displays the operational status of the interface, which can be Up (↑), Down (↓), or Unknown (?). The operational status of a network interface is determined by the status of its physical ports.

Administrative Status

Displays the administrative status of the interface, which can be Up (↑), Down (↓), or Unknown (?). You can control the administrative status of an interface when you make changes to the configuration or during maintenance. The administrative status can be different from the operational status. However, if the administrative status of a LIF is Down, the operational status is Down by default.

IP Address

Displays the IP address of the interface.

Role

Displays the role of the interface. Possible roles are Cluster-Management LIFs, Node-Management LIFs, Cluster LIFs, and Intercluster LIFs.

Home Port

Displays the physical port to which the interface was originally associated.

Current Port

Displays the physical port to which the interface is currently associated. After LIF migration, the current port might be different from the home port.

Failover Policy

Displays the failover policy that is configured for the interface.

Routing Groups

Displays the name of the routing group. You can view more information about the routes and the destination gateway by clicking the routing group name.

Routing groups are not supported for ONTAP 8.3 or later and therefore a blank column is displayed for these clusters.

Failover Group

Displays the name of the failover group.

Nodes tab

Displays information about nodes in the selected cluster. You can view detailed information about the HA pairs, disk shelves, and ports:

HA Details

Provides a pictorial representation of the HA state and the health status of the nodes in the HA pair. The health status of the node is indicated by the following colors:

- **Green**

The node is in a working condition.

- **Yellow**

The node has taken over the partner node or the node is facing some environmental issues.

- **Red**

The node is down.

You can view information about the availability of the HA pair and take required action to prevent any risks. For example, in the case of a possible takeover operation, the following message is displayed: Storage failover possible.

You can view a list of the events related to the HA pair and its environment, such as fans, power supplies, NVRAM battery, flash cards, service processor, and connectivity of disk shelves. You can also view the time when the events were triggered.

You can view other node-related information, such as the model number.

If there are single-node clusters, you can also view details about the nodes.

Disk Shelves

Displays information about the disk shelves in the HA pair.

You can also view events generated for the disk shelves and the environmental components, and the time when the events were triggered.

- **Shelf ID**

Displays the ID of the shelf where the disk is located.

- **Component Status**

Displays environmental details of the disk shelves, such as power supplies, fans, temperature sensors, current sensors, disk connectivity, and voltage sensors. The environmental details are displayed as icons in the following colors:

- **Green**

The environmental components are in working properly.

- **Grey**

No data is available for the environmental components.

- **Red**

Some of the environmental components are down.

- **State**

Displays the state of the disk shelf. The possible states are Offline, Online, No status, Initialization required, Missing, and Unknown.

- **Model**

Displays the model number of the disk shelf.

- **Local Disk Shelf**

Indicates whether the disk shelf is located on the local cluster or the remote cluster. This column is displayed only for clusters in a MetroCluster configuration.

- **Unique ID**

Displays the unique identifier of the disk shelf.

- **Firmware Version**

Displays the firmware version of the disk shelf.

Ports

Displays information about the associated FC, FCoE, and Ethernet ports. You can view details about the ports and the associated LIFs by clicking the port icons.

You can also view the events generated for the ports.

You can view the following port details:

- **Port ID**

Displays the name of the port. For example, the port names can be e0M, e0a, and e0b.

- **Role**

Displays the role of the port. The possible roles are Cluster, Data, Intercluster, Node-Management, and Undefined.

- **Type**

Displays the physical layer protocol used for the port. The possible types are Ethernet, Fibre Channel, and FCoE.

- WWPN

Displays the World Wide Port Name (WWPN) of the port.

- Firmware Rev

Displays the firmware revision of the FC/FCoE port.

- Status

Displays the current state of the port. The possible states are Up, Down, Link Not Connected, or Unknown (?).

You can view the port-related events from the Events list. You can also view the associated LIF details, such as LIF name, operational status, IP address or WWPN, protocols, name of the SVM associated with the LIF, current port, failover policy and failover group.

Disks tab

Displays details about the disks in the selected cluster. You can view disk-related information such as the number of used disks, spare disks, broken disks, and unassigned disks. You can also view other details such as the disk name, disk type, and the owner node of the disk.

Disk Pool Summary

Displays the number of disks, which are categorized by effective types (FCAL, SAS, SATA, MSATA, SSD, NVMe SSD, SSD CAP, Array LUN, and VMDISK), and the state of the disks. You can also view other details, such as the number of aggregates, shared disks, spare disks, broken disks, unassigned disks, and unsupported disks. If you click the effective disk type count link, disks of the selected state and effective type are displayed. For example, if you click the count link for the disk state Broken and effective type SAS, all disks with the disk state Broken and effective type SAS are displayed.

Disk

Displays the name of the disk.

RAID Groups

Displays the name of the RAID group.

Owner Node

Displays the name of the node to which the disk belongs. If the disk is unassigned, no value is displayed in this column.

State

Displays the state of the disk: Aggregate, Shared, Spare, Broken, Unassigned, Unsupported or Unknown. By default, this column is sorted to display the states in the following order: Broken, Unassigned, Unsupported, Spare, Aggregate, and Shared.

Local Disk

Displays either Yes or No to indicate whether the disk is located on the local cluster or the remote cluster. This

column is displayed only for clusters in a MetroCluster configuration.

Position

Displays the position of the disk based on its container type: for example, Copy, Data, or Parity. By default, this column is hidden.

Impacted Aggregates

Displays the number of aggregates that are impacted due to the failed disk. You can move the pointer over the count link to view the impacted aggregates and then click the aggregate name to view details of the aggregate. You can also click the aggregate count to view the list of impacted aggregates in the Health: All Aggregates view.

No value is displayed in this column for the following cases:

- For broken disks when a cluster containing such disks is added to Unified Manager
- When there are no failed disks

Storage Pool

Displays the name of the storage pool to which the SSD belongs. You can move the pointer over the storage pool name to view details of the storage pool.

Storable Capacity

Displays the disk capacity that is available for use.

Raw Capacity

Displays the capacity of the raw, unformatted disk before right-sizing and RAID configuration. By default, this column is hidden.

Type

Displays the types of disks: for example, ATA, SATA, FCAL, or VMDISK.

Effective Type

Displays the disk type assigned by ONTAP.

Certain ONTAP disk types are considered equivalent for the purposes of creating and adding to aggregates, and spare management. ONTAP assigns an effective disk type for each disk type.

Spare Blocks Consumed %

Displays in percentage the spare blocks that are consumed in the SSD disk. This column is blank for disks other than SSD disks.

Rated Life Used %

Displays in percentage an estimate of the SSD life used, based on the actual SSD usage and the manufacturer's prediction of SSD life. A value greater than 99 indicates that the estimated endurance has been consumed, but may not indicate SSD failure. If the value is unknown, then the disk is omitted.

Firmware

Displays the firmware version of the disk.

RPM

Displays the revolutions per minute (RPM) of the disk. By default, this column is hidden.

Model

Displays the model number of the disk. By default, this column is hidden.

Vendor

Displays the name of the disk vendor. By default, this column is hidden.

Shelf ID

Displays the ID of the shelf where the disk is located.

Bay

Displays the ID of the bay where the disk is located.

Related Annotations pane

Enables you to view the annotation details associated with the selected cluster. The details include the annotation name and the annotation values that are applied to the cluster. You can also remove manual annotations from the Related Annotations pane.

Related Devices pane

Enables you to view device details that are associated with the selected cluster.

The details include properties of the device that is connected to the cluster such as the device type, size, count, and health status. You can click on the count link for further analysis on that particular device.

You can use MetroCluster Partner pane to obtain count and also details on the remote MetroCluster partner along with its associated cluster components such as nodes, aggregates, and SVMs. The MetroCluster Partner pane is displayed only for clusters in a MetroCluster configuration.

The Related Devices pane enables you to view and navigate to the nodes, SVMs, and aggregates that are related to the cluster:

MetroCluster Partner

Displays the health status of the MetroCluster partner. Using the count link, you can navigate further and obtain information about the health and capacity of the cluster components.

Nodes

Displays the number, capacity, and health status of the nodes that belong to the selected cluster. Capacity indicates the total usable capacity over available capacity.

Storage Virtual Machines

Displays the number of SVMs that belong to the selected cluster.

Aggregates

Displays the number, capacity, and the health status of the aggregates that belong to the selected cluster.

Related Groups pane

Enables you to view the list of groups that includes the selected cluster.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts for the selected cluster. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Related information

[Volumes page](#) [Viewing the cluster list and details](#)

Aggregate / Health details page

You can use the Aggregate / Health details page to view detailed information about the selected aggregate, such as the capacity, disk information, configuration details, and events generated. You can also view information about the related objects and related alerts for that aggregate.

Command buttons



When monitoring a FabricPool-enabled aggregate, the committed and overcommitted values on this page are relevant only to the local, or performance tier, capacity. The amount of space available in the cloud tier is not reflected in the overcommitted values. Similarly, the aggregate threshold values are relevant only to the local performance tier.

The command buttons enable you to perform the following tasks for the selected aggregate:

- **Switch to Performance View**

Enables you to navigate to the Aggregate / Performance details page.

- **Actions**

- Add Alert

Enables you to add an alert to the selected aggregate.

- Edit Thresholds

Enables you to modify the threshold settings for the selected aggregate.

- **View Aggregates**

Enables you to navigate to the Health: All Aggregates view.

Capacity tab

The Capacity tab displays detailed information about the selected aggregate, such as its capacity, thresholds, and daily growth rate.

By default, capacity events are not generated for root aggregates. Also, the threshold values used by Unified Manager are not applicable to node root aggregates. Only a technical support representative can modify the settings for these events to be generated. When the settings are modified by a technical support representative, the threshold values are applied to the node root aggregate.

• Capacity

Displays the data capacity graph and the Snapshot copies graph, which display capacity details about the aggregate:

- Logical Space Used

The real size of the data that is being stored on the aggregate without applying the savings from using ONTAP storage efficiency technologies.

- Used

The physical capacity used by data in the aggregate.

- Overcommitted

When space in the aggregate is overcommitted, the chart displays a flag with the overcommitted amount.

- Warning

Displays a dotted line at the location where the warning threshold is set; meaning space in the aggregate is nearly full. If this threshold is breached, the Space Nearly Full event is generated.

- Error

Displays a solid line at the location where the error threshold is set; meaning space in the aggregate is full. If this threshold is breached, the Space Full event is generated.

- Snapshot Copies graph

This graph is displayed only when the used Snapshot capacity or the Snapshot reserve is not zero.

Both of the graphs display the capacity by which the Snapshot capacity exceeds the Snapshot reserve if the used Snapshot capacity exceeds the Snapshot reserve.

• Cloud Tier

Displays the space used by data in the cloud tier for FabricPool-enabled aggregates. A FabricPool can be either licensed or unlicensed.

When the cloud tier is mirrored to another cloud provider (the “mirror tier”) then both cloud tiers are displayed here.

• Details

Displays detailed information about capacity.

- Total Capacity

Displays the total capacity in the aggregate.

- Data Capacity

Displays the amount of space used by the aggregate (used capacity) and the amount of available space in the aggregate (free capacity).

- Snapshot Reserve

Displays the used and free Snapshot capacity of the aggregate.

- Overcommitted Capacity

Displays the aggregate overcommitment. Aggregate overcommitment enables you to provide more storage than is actually available from a given aggregate, as long as not all of that storage is currently being used. When thin provisioning is in use, the total size of volumes in the aggregate can exceed the total capacity of the aggregate.



If you have overcommitted your aggregate, you must monitor its available space carefully and add storage as required to avoid write errors due to insufficient space.

- Cloud Tier

Displays the space used by data in the cloud tier for FabricPool-enabled aggregates. A FabricPool can be either licensed or unlicensed. When the cloud tier is mirrored to another cloud provider (the mirror tier) then both cloud tiers are displayed here

- Total Cache Space

Displays the total space of the solid-state drives (SSDs) or allocation units that are added to a Flash Pool aggregate. If you have enabled Flash Pool for an aggregate but have not added any SSDs, then the cache space is displayed as 0 KB.



This field is hidden if Flash Pool is disabled for an aggregate.

- Aggregate Thresholds

Displays the following aggregate capacity thresholds:

- Nearly Full Threshold

Specifies the percentage at which an aggregate is nearly full.

- Full Threshold

Specifies the percentage at which an aggregate is full.

- Nearly Overcommitted Threshold

Specifies the percentage at which an aggregate is nearly overcommitted.

- Overcommitted Threshold

Specifies the percentage at which an aggregate is overcommitted.

- Other Details: Daily Growth Rate

Displays the disk space used in the aggregate if the rate of change between the last two samples continues for 24 hours.

For example, if an aggregate uses 10 GB of disk space at 2 pm and 12 GB at 6 pm, the daily growth rate (GB) for this aggregate is 2 GB.

- Volume Move

Displays the number of volume move operations that are currently in progress:

- Volumes Out

Displays the number and capacity of the volumes that are being moved out of the aggregate.

You can click the link to view more details, such as the volume name, aggregate to which the volume is moved, status of the volume move operation, and the estimated end time.

- Volumes In

Displays the number and remaining capacity of the volumes that are being moved into the aggregate.

You can click the link to view more details, such as the volume name, aggregate from which the volume is moved, status of the volume move operation, and the estimated end time.

- Estimated used capacity after volume move

Displays the estimated amount of used space (as a percentage, and in KB, MB, GB, and so on) in the aggregate after the volume move operations are complete.

- **Capacity Overview - Volumes**

Displays graphs that provide information about the capacity of the volumes contained in the aggregate. The amount of space used by the volume (used capacity) and the amount of available space (free capacity) in the volume is displayed. When the Thin-Provisioned Volume Space At Risk event is generated for thinly provisioned volumes, the amount of space used by the volume (used capacity) and the amount of space that is available in the volume but cannot be used (unusable capacity) because of aggregate capacity issues is displayed.

You can select the graph you want to view from the drop-down lists. You can sort the data displayed in the graph to display details such as the used size, provisioned size, available capacity, fastest daily growth rate, and slowest growth rate. You can filter the data based on the storage virtual machines (SVMs) that contain the volumes in the aggregate. You can also view details for thinly provisioned volumes. You can view the details of specific points on the graph by positioning your cursor over the area of interest. By default, the graph displays the top 30 filtered volumes in the aggregate.

Disk Information tab

Displays detailed information about the disks in the selected aggregate, including the RAID type and size, and

the type of disks used in the aggregate. The tab also graphically displays the RAID groups, and the types of disks used (such as SAS, ATA, FCAL, SSD, or VMDISK). You can view more information, such as the disk's bay, shelf, and rotational speed, by positioning your cursor over the parity disks and data disks.

- **Data**

Graphically displays details about dedicated data disks, shared data disks, or both. When the data disks contain shared disks, graphical details of the shared disks are displayed. When the data disks contain dedicated disks and shared disks, graphical details of both the dedicated data disks and the shared data disks are displayed.

- **RAID Details**

RAID details are displayed only for dedicated disks.

- **Type**

Displays the RAID type (RAID0, RAID4, RAID-DP, or RAID-TEC).

- **Group Size**

Displays the maximum number of disks allowed in the RAID group.

- **Groups**

Displays the number of RAID groups in the aggregate.

- **Disks Used**

- **Effective Type**

Displays the types of data disks (for example, ATA, SATA, FCAL, SSD, or VMDISK) in the aggregate.

- **Data Disks**

Displays the number and capacity of the data disks that are assigned to an aggregate. Data disk details are not displayed when the aggregate contains only shared disks.

- **Parity Disks**

Displays the number and capacity of the parity disks that are assigned to an aggregate. Parity disk details are not displayed when the aggregate contains only shared disks.

- **Shared Disks**

Displays the number and capacity of the shared data disks that are assigned to an aggregate. Shared disk details are displayed only when the aggregate contains shared disks.

- **Spare Disks**

Displays the disk effective type, number, and capacity of the spare data disks that are available for the node in the selected aggregate.



When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

- **SSD Cache**

Provides details about dedicated cache SSD disks and shared cache SSD disks.

The following details for the dedicated cache SSD disks are displayed:

- **RAID Details**

- **Type**

Displays the RAID type (RAID0, RAID4, RAID-DP or RAID-TEC).

- **Group Size**

Displays the maximum number of disks allowed in the RAID group.

- **Groups**

Displays the number of RAID groups in the aggregate.

- **Disks Used**

- **Effective Type**

Indicates that the disks used for cache in the aggregate are of type SSD.

- **Data Disks**

Displays the number and capacity of the data disks that are assigned to an aggregate for cache.

- **Parity Disks**

Displays the number and capacity of the parity disks that are assigned to an aggregate for cache.

- **Spare Disks**

Displays the disk effective type, number, and capacity of the spare disks that are available for the node in the selected aggregate for cache.



When an aggregate is failed over to the partner node, Unified Manager does not display all of the spare disks that are compatible with the aggregate.

Provides the following details for the shared cache:

- **Storage Pool**

Displays the name of the storage pool. You can move the pointer over the storage pool name to view the following details:

- **Status**

Displays the status of the storage pool, which can be healthy or unhealthy.

- **Total Allocations**

Displays the total allocation units and the size in the storage pool.

- **Allocation Unit Size**

Displays the minimum amount of space in the storage pool that can be allocated to an aggregate.

- **Disks**

Displays the number of disks used to create the storage pool. If the disk count in the storage pool column and the number of disks displayed in the Disk Information tab for that storage pool do not match, then it indicates that one or more disks are broken and the storage pool is unhealthy.

- **Used Allocation**

Displays the number and size of the allocation units used by the aggregates. You can click the aggregate name to view the aggregate details.

- **Available Allocation**

Displays the number and size of the allocation units available for the nodes. You can click the node name to view the aggregate details.

- **Allocated Cache**

Displays the size of the allocation units used by the aggregate.

- **Allocation Units**

Displays the number of allocation units used by the aggregate.

- **Disks**

Displays the number of disks contained in the storage pool.

- **Details**

- **Storage Pool**

- Displays the number of storage pools.

- **Total Size**

- Displays the total size of the storage pools.

- **Cloud Tier**

Displays the name of the cloud tier, if you have configured a FabricPool-enabled aggregate, and shows the total space used. When the cloud tier is mirrored to another cloud provider (the mirror tier) then the details for both cloud tiers are displayed here

Configuration tab

The Configuration tab displays details about the selected aggregate, such as its cluster node, block type, RAID type, RAID size, and RAID group count:

- **Overview**

- **Node**

Displays the name of the node that contains the selected aggregate.

- Block Type

Displays the block format of the aggregate: either 32-bit or 64-bit.

- RAID Type

Displays the RAID type (RAID0, RAID4, RAID-DP, RAID-TEC or Mixed RAID).

- RAID Size

Displays the size of the RAID group.

- RAID Groups

Displays the number of RAID groups in the aggregate.

- SnapLock Type

Displays the SnapLock Type of the aggregate.

- **Cloud Tier**

If this is a FabricPool-enabled aggregate, the details for the cloud tier are displayed. Some fields are different depending on the storage provider. When the cloud tier is mirrored to another cloud provider (the “mirror tier”) then both cloud tiers are displayed here.

- Provider

Displays the name of the storage provider, for example, StorageGRID, Amazon S3, IBM Cloud Object Storage, Microsoft Azure Cloud, Google Cloud Storage, or Alibaba Cloud Object Storage.

- Name

Displays the name of the cloud tier when it was created by ONTAP.

- Server

Displays the FQDN of the cloud tier.

- Port

The port being used to communicate with the cloud provider.

- Access Key or Account

Displays the access key or account for the cloud tier.

- Container Name

Displays the bucket or container name of the cloud tier.

- SSL

Displays whether SSL encryption is enabled for the cloud tier.

History area

The History area displays graphs that provide information about the capacity of the selected aggregate. Additionally, you can click the **Export** button to create a report in CSV format for the chart that you are viewing.

You can select a graph type from the drop-down list at the top of the History pane. You can also view details for a specific time period by selecting either 1 week, 1 month, or 1 year. History graphs can help you identify trends: for example, if the aggregate usage is consistently breaching the Nearly Full threshold, you can take the appropriate action.

History graphs display the following information:

- **Aggregate Capacity Used (%)**

Displays the used capacity in the aggregate and the trend in how aggregate capacity is used based on the usage history as line graphs, in percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Capacity Used legend, the Capacity Used graph line is hidden.

- **Aggregate Capacity Used vs Total Capacity**

Displays the trend in how aggregate capacity is used based on the usage history, as well as the used capacity and the total capacity, as line graphs, in bytes, kilobytes, megabytes, and so on, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Trend Capacity Used legend, the Trend Capacity Used graph line is hidden.

- **Aggregate Capacity Used (%) vs Committed (%)**

Displays the trend in how aggregate capacity is used based on the usage history, as well as the committed space as line graphs, as a percentage, on the vertical (y) axis. The time period is displayed on the horizontal (x) axis. You can select a time period of a week, a month, or a year. You can view the details for specific points on the graph by positioning your cursor over a particular area. You can hide or display a line graph by clicking the appropriate legend. For example, when you click the Space Committed legend, the Space Committed graph line is hidden.

Events list

The Events list displays details about new and acknowledged events:

- **Severity**

Displays the severity of the event.

- **Event**

Displays the event name.

- **Triggered Time**

Displays the time that has elapsed since the event was generated. If the time elapsed exceeds a week, the timestamp for when the event was generated is displayed.

Related Devices pane

The Related Devices pane enables you to view the cluster node, volumes, and disks that are related to the aggregate:

- **Node**

Displays the capacity and the health status of the node that contains the aggregate. Capacity indicates the total usable capacity over available capacity.

- **Aggregates in the Node**

Displays the number and capacity of all the aggregates in the cluster node that contains the selected aggregate. The health status of the aggregates is also displayed, based on the highest severity level. For example, if a cluster node contains ten aggregates, five of which display the Warning status and the remaining five of which display the Critical status, then the status displayed is Critical.

- **Volumes**

Displays the number and capacity of FlexVol volumes and FlexGroup volumes in the aggregate; the number does not include FlexGroup constituents. The health status of the volumes is also displayed, based on the highest severity level.

- **Resource Pool**

Displays the resource pools related to the aggregate.

- **Disks**

Displays the number of disks in the selected aggregate.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected aggregate. You can also add an alert by clicking the Add Alert link or edit an existing alert by clicking the alert name.

Related information

[Viewing storage pool details](#)

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.