



Managing and monitoring clusters and cluster object health

Active IQ Unified Manager 9.7

NetApp
June 10, 2024

Table of Contents

- Managing and monitoring clusters and cluster object health 1
 - Understanding cluster monitoring 1
 - Viewing the cluster list and details 2
 - Checking the health of clusters in a MetroCluster configuration 3
 - Viewing the health and capacity status of All SAN Array clusters 4
 - Viewing the node list and details 5
 - Generating a hardware inventory report for contract renewal 6
 - Viewing the Storage VM list and details 6
 - Viewing the aggregate list and details 7
 - Viewing FabricPool capacity information 7
 - Viewing storage pool details 9
 - Viewing the volume list and details 9
 - Viewing details about NFS shares 10
 - Viewing details about SMB/CIFS shares 10
 - Viewing the list of Snapshot copies 11
 - Deleting Snapshot copies 12
 - Calculating reclaimable space for Snapshot copies 12
 - Description of cluster object windows and dialog boxes 13

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.

Before you begin

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

About this task

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.

Checking the health of clusters in a MetroCluster configuration

You can use Unified Manager to check the operational health of clusters, and their components, in a MetroCluster configuration. 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.

Before you begin

- 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 must be monitored by the same instance of Unified Manager.

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.
3. Next to a MetroCluster event, click the name of the cluster.

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



If no MetroCluster events are displayed, you can use the Search bar to search for the name of the cluster involved in the performance 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.

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.

Before you begin

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

About this task

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.

Before you begin

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

About this task

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.

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.

Before you begin

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

Steps

1. In the left navigation pane, click **Storage > Nodes**.
2. In the **Health: All Nodes** view, or the **Performance: All Nodes** view, click the **Hardware Inventory Report** button.

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

3. 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.

Before you begin

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

About this task

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.

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.

Before you begin

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

About this task

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.

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.

About this task

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>
Aggregates	<ol style="list-style-type: none"> a. On the Capacity: All Aggregates view, click an aggregate where the Type field indicates “SSD (FabricPool)”. b. On the Aggregate / Health details page, click the Capacity tab. <p>The display shows the total capacity used in the cloud tier.</p> <ol style="list-style-type: none"> c. Click the Disk Information tab. <p>The display shows the name of the cloud tier and the capacity used.</p> <ol style="list-style-type: none"> d. Click the Configuration tab. <p>The display shows the name of the cloud tier and other detailed information about the cloud tier.</p>
Volumes	<ol style="list-style-type: none"> a. On the Capacity: All Volumes view, click a volume where a policy name appears in the “Tiering Policy” field. b. On the Volume / Health details page, click the Configuration tab. <p>The display shows the name of the FabricPool tiering policy assigned to the volume.</p>

After you finish

For more information on FabricPool aggregates, see the *ONTAP 9 Disks and Aggregates Power Guide*.

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.

Before you begin

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

Steps

1. In the left navigation pane, click **Storage > Aggregates**.
2. 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.

Before you begin

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

About this task

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

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).

Before you begin

- 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.

Steps

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.

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).

Before you begin

- 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.

About this task



Shares in folders are not displayed.

Steps

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.

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.

Before you begin

- 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.

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.

Before you begin

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.

About this task

- 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.

Before you begin

- 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.

Health: All Clusters view

The Health: All Clusters view enables you to view health information about the clusters that you are monitoring.

By default, objects in the view pages are sorted based on event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

You can associate a cluster to a predefined annotation by using the **Annotate** button.

See [Cluster health fields](#) for descriptions of all the fields on this page.

Cluster health fields

The following fields are available in the Health: All Clusters view and can be used in custom views and in reports.

- **Status**

An icon that identifies the current status of the cluster. The status can be Critical (❌), Error (⚠️), Warning (⚠️), or Normal (✅).

- **Cluster**

The name of the cluster. You can click the cluster name to navigate to that cluster's health details page.

- **Cluster FQDN**

The fully qualified domain name (FQDN) of the cluster.

- **Communication Status**

Whether the cluster is reachable or not.

The status is displayed as Good if the cluster is reachable. If the cluster is not reachable or if the login credentials are invalid, the status is displayed as Not Reachable.

- **System Health**

High-level information about the status of the cluster, which is calculated based on the status of various cluster subsystems.

Possible values are OK, OK with suppressed, Degraded, and Components not reachable. These values are determined by the health monitors in ONTAP software.

- **Last Refreshed Time**

The timestamp of when the monitoring samples of the cluster were last collected.

- **FIPS Enabled**

Whether FIPS mode is enabled on the cluster.

- **OS Version**

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.

- **Node Count**

The number of nodes that belong to the cluster.

- **Host Name or IP Address**

The FQDN, short name, or the IP address of the cluster-management LIF that is used to connect to the cluster.

- **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.

- **Personality**

Identifies if this is an All SAN Array configured cluster.

- **Serial Number**

The serial number of the cluster.

- **Contact**

The contact information of the cluster.

- **Location**

The location of the cluster.

Capacity: All Clusters view

The Capacity: All Clusters view enables you to view summarized information about storage capacity and utilization in all clusters. This information helps you to understand possible capacity risks and to take appropriate action to rebalance workloads.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

See [Cluster capacity fields](#) for descriptions of all the fields on this page.

Cluster capacity fields

The following fields are available in the Capacity: All Clusters view and can be used in custom views and in reports.

- **Cluster**

The cluster name. You can click the cluster name to navigate to that cluster's capacity details page.

- **Cluster FQDN**

The fully qualified domain name (FQDN) of the cluster.

- **HA Pair**

The HA pair value obtained by forming two nodes.

- **Total Raw Capacity**

Displays the total physical capacity of all disks in the array.

- **Unconfigured Raw Capacity**

The unconfigured capacity of disks whose container type is other than aggregate, broken, spare, or

shared. This capacity is always higher than the physical capacity of the disk in ONTAP. For example, consider a 2 TB disk. The physical capacity of the disk is 1.6 TB in ONTAP whereas the unconfigured raw capacity in Unified Manager is 1.8 TB.

- **Aggregate Total Capacity**

The total size of the available aggregates for the user. This includes the Snapshot copy reserve.

- **Aggregate Used Capacity**

The capacity already in use on aggregates. This includes the capacity consumed by volumes, LUNs, and other storage efficiency technology overheads.

- **Aggregate Unused Capacity**

The capacity that might be available for storing additional data on the aggregate. This includes the Snapshot copy reserve.

- **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.

- **Allocated LUN Capacity**

The capacity of LUNs that are mapped.

- **Unallocated LUN Capacity**

The capacity of all LUNs not mapped to the Host.

- **Volume Total Capacity**

The total capacity of the volumes (used plus unused).

- **Volume Used Capacity**

The used capacity of the volumes.

- **Volume Unused Capacity**

The unused capacity of the volumes.

- **Volume Protection Capacity**

The capacity of volumes that have SnapMirror and SnapVault enabled.

- **Cluster Licensed Cloud Tier Used**

The space used by data in the cloud tier for storage providers that require a FabricPool license.

- **Cluster StorageGRID Capacity Used**

The space used by data in the cloud tier for StorageGRID systems that do not require a FabricPool license.

- **Model/Family**

The model or family name of the cluster.

- **OS Version**

The version of ONTAP installed on the system.

- **Contact**

The contact information of the cluster.

- **Location**

The location of the cluster.

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.

- 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.

- 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 capacity used for all connected cloud tiers 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, serial number, 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.

- Serial Number

Displays the serial number of the cluster.

- 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.

- Serial Number

Displays the serial number of the remote cluster.

- Location

Displays the location of the remote cluster.

• **MetroCluster Overview**

Provides details about the local cluster in a MetroCluster configuration. This information is displayed only for MetroCluster configurations.

- Type

Displays whether the MetroCluster type is two-node or four-node.

- Configuration

Displays the MetroCluster configuration, which can have the following values:

- 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.

- 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.

• **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.

- **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 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 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. 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 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 and the serial 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, 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.

Health: All Nodes view

The Health: All Nodes view enables you to view detailed information about the nodes in all clusters being managed by Unified Manager.

By default, objects in the view pages are sorted based on event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.





The **Hardware Inventory Report** button is provided when Unified Manager, and the clusters it is managing, are installed in a site with no external network connectivity. This button generates a `.csv` file 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 reporting functionality is helpful for contract renewal within secure sites that are not connected to the NetAppActive IQ platform.

See [Node health fields](#) for descriptions of all the fields on this page.

Node health fields

The following fields are available in the Health: All Nodes view and can be used in custom views and in reports.

- **Status**

An icon that identifies the current status of the node. The status can be Critical () , Error () , Warning () , or Normal () .

- **Node**

The name of the node. You can click the node name to navigate to that cluster's node details page.

- **State**

The state of the node. The state can be Up or Down.

- **HA State**

The state of the HA pair. The state can be Error, Warning, Normal, or Not applicable.

- **Down Time**

The time that has elapsed or the timestamp since the node is offline. If the time elapsed exceeds a week, the timestamp when the node went offline is displayed.

- **All Flash Optimized**

Whether the node is optimized to support only solid-state drives (SSDs).

- **Model/Family**

The model of the node.

- **OS version**

The ONTAP software version that the node is running.

- **Serial Number**

The serial number of the node.

- **Firmware Version**

The firmware version number of the node.

- **Aggregate Used Capacity**

The amount of space used for data in the node's aggregates.

- **Aggregate Total Capacity**

The total space available for data in the node's aggregates.

- **Usable Spare Capacity**

The amount of available space in the node that can be used to enhance the aggregate capacity.

- **Usable Raw Capacity**

The amount of space that is usable in the node.

- **Total Raw Capacity**

The capacity of every unformatted disk in the node before right-sizing and RAID configuration.

- **Storage VM Count**

The number of SVMs contained by the cluster.

- **FC Port Count**

The number of FC ports contained by the node.

- **FCoE Port Count**

The number of FCoE ports contained by the node.

- **Ethernet Port Count**

The number of ethernet ports contained by the node.

- **Flash Card Size**

The size of the flash cards installed on the node.

- **Flash Card Count**

The number of flash cards installed on the node.

- **Disk Shelves Count**

The number of disk shelves contained by the node.

- **Disk Count**

The number of disks in the node.

- **Cluster**

The name of the cluster to which the node belongs. You can click the cluster name to navigate to that cluster's health details page.

- **Cluster FQDN**

The fully qualified domain name (FQDN) of the cluster.

- **Contact**

The contact information of the node.

- **Location**

The location of the node.

Health: All Aggregates view

The Health: All Aggregates view displays information about the aggregates that are monitored, and enables you to view and modify the threshold settings.

By default, objects in the view pages are sorted based on event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

You can customize capacity threshold settings for one or more aggregates by using the **Edit Threshold** button.

See [Aggregate health fields](#) for descriptions of all the fields on this page.

Aggregate health fields

The following fields are available in the Health: All Aggregates view and can be used in custom views and in reports.

- **Status**

The current status of the aggregate. The status can be Critical (❌), Error (⚠️), Warning (⚠️), or Normal (✅).

- **Aggregate**

The name of the aggregate.

- **State**

The current state of the aggregate:

- Offline

Read or write access is not allowed.

- Online

Read and write access to volumes hosted on this aggregate is allowed.

- Restricted

Limited operations, such as parity reconstruction, are allowed, but data access is not allowed.

- Creating
The aggregate is being created.
- Destroying
The aggregate is being destroyed.
- Failed
The aggregate cannot be brought online.
- Frozen
The aggregate is (temporarily) not serving requests.
- Inconsistent
The aggregate has been marked corrupted; contact technical support.
- Iron Restricted
Diagnostic tools cannot be run on the aggregate.
- Mounting
The aggregate is being mounted.
- Partial
At least one disk was found for the aggregate, but two or more disks are missing.
- Quiesced
The aggregate is quiesced.
- Quiescing
The aggregate is being quiesced.
- Reverted
The revert operation of the aggregate is completed.
- Unmounted
The aggregate is offline.
- Unmounting
The aggregate is being taken offline.
- Unknown
Specifies that the aggregate is discovered, but the aggregate information is not yet retrieved by the Unified Manager server.

- **Node**

The name of the node that contains the aggregate.

- **Mirror Status**

The mirror status of the aggregate:

- Mirrored

The aggregate plex data is mirrored.

- Mirror degraded

The aggregate plex data cannot be mirrored.

- Mirror resynchronizing

The aggregate plex data is being mirrored.

- Failed

The aggregate plex data mirroring failed.

- Invalid configuration

The initial state before an aggregate is created.

- Uninitialized

The aggregate is being created.

- Unmirrored

The aggregate is not mirrored.

- CP count check in progress

The aggregate has been assimilated and Unified Manager is validating that the CP counts for the plexes is similar.

- Limbo

There is an issue with the aggregate labels. The ONTAP system identifies the aggregate but cannot accurately assimilate the aggregate.

- Needs CP count check

The aggregate is assimilated but the CP counts on both plexes are not yet validated to be similar.

When an aggregate is in the mirror_resynchronizing state, then the resynchronization percentage is also shown.

- **In Transition**

Whether the aggregate has completed transition or not.

- **Type**

The aggregate type:

- HDD

- Hybrid

Combines HDDs and SSDs, but Flash Pool has not been enabled.

- Hybrid (Flash Pool)

Combines HDDs and SSDs, and Flash Pool has been enabled.

- SSD

- SSD (FabricPool)

Combines SSDs and a cloud tier

- VMDisk (SDS)

Virtual disks within a virtual machine

- VMDisk (FabricPool)

Combines virtual disks and a cloud tier

- LUN (FlexArray)

- **SnapLock Type**

The aggregate SnapLock Type. The possible values are Compliance, Enterprise, Non-SnapLock.

- **Used Data %**

The percentage of space used for data in the aggregate.

- **Available Data %**

The percentage of space available for data in the aggregate.

- **Used Data Capacity**

The amount of space used for data in the aggregate.

- **Available Data Capacity**

The amount of space available for data in the aggregate.

- **Total Data Capacity**

The total data size of the aggregate.

- **Committed Capacity**

The total space committed for all of the volumes in the aggregate.

When Autogrow is enabled on volumes that reside on the aggregate, the committed capacity is based on the maximum volume size set by autogrow, not based on the original volume size. For FabricPool aggregates, this value is relevant only to the local, or performance tier, capacity. The amount of space available in the cloud tier is not reflected in this value.

- **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.

- **Space Savings**

The storage efficiency ratio based on the total logical space that is being used to store the data and the total physical space that would be required to store the data without using ONTAP storage efficiency technologies.

This field is populated only when the monitored storage system is running ONTAP version 9.0 or greater, and only for non-root aggregates.

- **Cloud Tier Space Used**

The amount of space being used in the cloud tier; if the aggregate is a FabricPool aggregate.

- **RAID Type**

The RAID configuration type:

- RAID 0: All the RAID groups are of type RAID 0.
- RAID 4: All the RAID groups are of type RAID 4.
- RAID-DP: All the RAID groups are of type RAID-DP.
- RAID-TEC: All the RAID groups are of type RAID-TEC.
- Mixed RAID: The aggregate contains RAID groups of different RAID types (RAID 0, RAID 4, RAID-DP, and RAID-TEC).

- **Cluster**

The name of the cluster on which the aggregate resides. You can click the cluster name to navigate to that cluster's health details page.

- **Cluster FQDN**

The fully qualified domain name (FQDN) of the cluster.

Capacity: All Aggregates view

The Capacity: All Aggregates view enables you to view information about the capacity and utilization of aggregates in all clusters. This information enables you to understand possible capacity risks and also to view the configured, used, and unused capacity of aggregates.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed

data, add/remove/reorder columns on the page, and export the data on the page to a .csv or .pdf document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

See [Aggregate capacity fields](#) for descriptions of all the fields on this page.

Aggregate capacity fields

The following fields are available in the Aggregate Capacity and Utilization report and can be used in custom views and in reports.

- **Aggregate**

The aggregate name.

- **Daily Growth Rate %**

The growth rate that occurs every 24 hours in the aggregate.

- **Days To Full**

The estimated number of days remaining before the aggregate reaches full capacity.

- **Overcommitted Capacity %**

The aggregate overcommitment as a percentage.

- **Available Data %**

The available data capacity as a percentage.

- **Available Data Capacity**

The available data capacity.

- **Used Data %**

The used data capacity as a percentage.

- **Used Data Capacity**

The used data capacity.

- **Total Data Capacity**

The total data capacity (used plus available).

- **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.

- **Snapshot Reserve Available %**

The amount of space available for Snapshot copies as a percentage.

- **Snapshot Reserve Available Capacity**

The amount of space available for Snapshot copies.

- **Snapshot Reserve Used %**

The amount of space used by Snapshot copies from the snapshot reserve as a percentage.

- **Snapshot Reserve Used Capacity**

The amount of space used by snapshot copies from the snapshot reserve.

- **Snapshot Reserve Total Capacity**

The total snapshot reserve capacity of the aggregate.

- **Cloud Tier Space Used**

The amount of data capacity that is currently being used in the cloud tier.

- **Cloud Tier**

The name of the cloud tier object store when it was created by ONTAP.

- **State**

The current state of the aggregate.

- **Type**

The aggregate type:

- HDD

- Hybrid

Combines HDDs and SSDs, but Flash Pool has not been enabled.

- Hybrid (Flash Pool)

Combines HDDs and SSDs, and Flash Pool has been enabled.

- SSD

- SSD (FabricPool)

Combines SSDs and a cloud tier

- VMDisk (SDS)

Virtual disks within a virtual machine

- VMDisk (FabricPool)

Combines virtual disks and a cloud tier

- LUN (FlexArray)

- **RAID Type**

The RAID configuration type.

- **SnapLock Type**

The aggregate SnapLock Type. The possible values are Compliance, Enterprise, Non-SnapLock.

- **HA Pair**

The HA pair value obtained by forming two nodes.

- **Node**

The name of the node that contains the aggregate.

- **Cluster**

The cluster name. You can click the cluster name to navigate to that cluster's capacity details page.

- **Cluster FQDN**

The fully qualified domain name (FQDN) of the cluster.

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.

Health: All Storage VMs view

The Health: All Storage VMs view enables you to view detailed information about the storage virtual machines (SVMs) that you are monitoring.

By default, objects in the view pages are sorted based on event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a .csv or .pdf document. After you have customized the page, you can save the results as a custom view and then schedule a report of

this data to be generated and emailed on a regular basis.

You can associate an SVM to a predefined annotation by using the **Annotate** button.

See [SVM health fields](#) for descriptions of all the fields on this page.

SVM health fields

The following fields are available in the Health: All Storage VMs view and can be used in custom views and in reports.

- **Status**

The current status of the SVM. The status can be Critical (❌), Error (⚠️), Warning (⚠️), or Normal (✅).

- **Storage VM**

The name of the SVM.

- **State**

The current administrative state of the SVM. The state can be Running, Stopped, Starting, or Stopping.

- **Allowed Volume Type**

The type of volume that can be created in the SVM. The type can be FlexVol or FlexVol/FlexGroup.

- **Allowed Protocols**

The type of protocols that can be configured on the SVM. The available protocols are FC/FCoE, iSCSI, HTTP, NDMP, NVMe, NFS, and CIFS.

- **Available Data Capacity**

The available data capacity of all the volumes in the SVM.

- **Total Data Capacity**

The total data capacity of all the volumes in the SVM.

- **Root Volume**

The name of the root volume of the SVM.

- **NIS State**

The state of the Network Information Service (NIS). The state can be Enabled, Disabled, or Not Configured.

- **NIS Domain**

The NIS domain name. This column is blank when the NIS server is disabled or is not configured.

- **DNS State**

The state of the Domain Name System (DNS). The state can be Enabled, Disabled, or Not Configured.

- **DNS Domain**

The DNS domain name.

- **Name Service Switch**

The information type gathered from hosts. Possible values are file, LDAP, or NIS.

- **LDAP Enabled**

Whether the LDAP protocol is enabled or not.

- **Maximum Allowed Volumes**

The maximum allowed volumes that can be configured on the SVM.

- **Volume Count**

The number of volumes contained by the SVM.

- **Cluster**

The name of the cluster to which the SVM belongs.

- **Cluster FQDN**

The fully qualified domain name (FQDN) of the cluster.

Storage VM / Health details page

You can use the Storage VM / Health details page to view detailed information about the selected SVM, such as its health, capacity, configuration, data policies, logical interfaces (LIFs), LUNs, qtrees, and user and user group quotas. You can also view information about the related objects and related alerts for the SVM.



You can monitor only data SVMs.

Command buttons

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

- **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 SVM.

- Annotate

Enables you to annotate the selected SVM.

- **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 SVM. 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 SVM. For example, information is displayed about aggregates that are likely to breach the set threshold values.

- **Protection Issues**

Provides a quick overview of SVM protection-related health by displaying, as a graph, the total number of relationships, including relationships that have protection issues and relationships that do not have any protection-related issues. 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 SVM. 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 SVM. For example, information is displayed about volumes that have a Snapshot copy reserve that is almost full or about SnapMirror relationship lag issues.

If the selected SVM is a repository SVM, the Protection area does not display.

Capacity tab

The Capacity tab displays detailed information about the data capacity of the selected SVM.

The following information is displayed for an SVM 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 SVM.

- Used

Displays the space used by data in the volumes that belong to the SVM.

- Guaranteed Available

Displays the guaranteed available space for data that is available for volumes in the SVM.

- Unguaranteed

Displays the available space remaining for data that is allocated for thinly provisioned volumes in the SVM.

- **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 SVM, such as its cluster, root volume, the type of volumes it contains (FlexVol volumes), and the policies created on the SVM:

- **Overview**

- Cluster

Displays the name of the cluster to which the SVM belongs.

- Allowed Volume Type

Displays the type of volumes that can be created in the SVM. The type can be FlexVol or FlexVol/FlexGroup.

- Root Volume

Displays the name of the root volume of the SVM.

- Allowed Protocols

Displays the type of protocols that can be configured on the SVM. 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 SVM. Also, indicates if the interfaces are up (●) or down (●).

- SAN

Displays the number of SAN interfaces that are associated with the SVM. Also, indicates if the interfaces are up (●) or down (●).

- FC-NVMe

Displays the number of FC-NVMe interfaces that are associated with the SVM. Also, indicates if the interfaces are up (●) or down (●).

- **Management Network Interfaces**

- Availability

Displays the number of management interfaces that are associated with the SVM. Also, indicates if the management interfaces are up (●) or down (●).

- **Policies**

- Snapshots

Displays the name of the Snapshot policy that is created on the SVM.

- 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.

- **Services**

- Type

Displays the type of service that is configured on the SVM. 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 SVM:

- **Network Interface**

Displays the name of the interface that is created on the selected SVM.

- **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 SVM, for all SVMs in the current cluster, or for all SVMs 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 SVM. 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 SVM. 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 SVM. You can view information such as the status of the SMB share, share name, path associated with the SVM, 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 SVM.

- **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 SVM, 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 SVM.

- **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 SVM. 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 SVM. 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. You can also view information whether thin provisioning is enabled on the LUN and if the LUN is mapped to an initiator group.

You can also 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 SVM.

You can also view the LUNs and 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 SVM. Details include the annotation name and the annotation values that are applied to the SVM. 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 SVM:

- **Cluster**

Displays the health status of the cluster to which the SVM belongs.

- **Aggregates**

Displays the number of aggregates that belong to the selected SVM. The health status of the aggregates is also displayed, based on the highest severity level. For example, if an SVM 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 SVM. 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 SVM. The health status of the volumes is also displayed, based on the highest severity level. When there are FlexGroup volumes in the SVM, 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 SVM.

Related Alerts pane

The Related Alerts pane enables you to view the list of alerts that are created for the selected SVM. You can also add an alert by clicking the **Add Alert** link or edit an existing alert by clicking the alert name.

Storage Pool dialog box

The Storage Pool dialog box enables you to view the details of the dedicated cache of SSDs, also known as *storage pools*. You can monitor the storage pools and view details such as the storage pool health, total and available cache, and used and available allocations in the storage pool.

You can view the following storage pool 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.

- **Cache Allocations**

- Used Allocations

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 Allocations


Displays the number and size of the allocation units available for the nodes. You can click the node name to view the aggregate details.

Health: All Volumes view

The Health: All Volumes view displays information about the volumes in the storage systems that are monitored and enables you to modify the volume threshold settings.

By default, objects in the view pages are sorted based on event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues that must be addressed.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

If you want to analyze the latency and throughput of a specific object, click the more icon  , then **Analyze Workload** and you can view performance and capacity charts on the Workload Analysis page.

See [Volume health fields](#) for descriptions of all the fields on this page.

Command buttons

- **Edit Threshold**

Displays the Edit Thresholds dialog box, which enables you to edit the health threshold settings for one or more volumes.

- **Protect**

Displays the following submenus:

- SnapMirror

Enables you to create a SnapMirror relationship for the selected volumes.

- SnapVault

Enables you to create a SnapVault relationship for the selected volumes.

- **Restore**

Displays the Restore dialog box, which enables you to restore directories or files from one volume at a time.

- **Annotate**

Enables you to annotate the selected volume.

Volume health fields

The following fields are available in the Health: All Volumes view and can be used in custom views and in reports.

- **Status**

The current status of a volume. The status can be Critical (❌), Error (⚠️), Warning (⚠️), or Normal (✅).

- **Volume**

The name of the volume.

- **Storage VM**

The SVM that contains the volume.

- **State**

The current state of the volume:

- 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.

- **Protection Role**

The protection role of a volume:

- Unprotected

A read/write volume with no outgoing or incoming SnapMirror or SnapVault relationships

- Protected

A read/write volume with an outgoing SnapMirror or SnapVault relationship

- Destination

A data protection (DP) volume or read/write volume with an incoming SnapMirror or SnapVault relationship

- Not Applicable

A volume for which protection roles do not apply, such as a load sharing volume, data constituent, or temporary volume

Clicking the role displays the Protection tab of the Volume / Health details page.

- **Style**

The style of volume; FlexVol or FlexGroup.

- **Mount Path**

The path to which the volume is mounted.

- **Available Data %**

The percentage of physical space currently available for data in the volume.

- **Available Data Capacity**

The amount of physical space currently available for data in the volume.

- **Used Data %**

The percentage of physical space used by data in the volume based on the total available data capacity.

- **Used Data Capacity**

The amount of physical space used by data in the volume.

- **Total Data Capacity**

The total physical space available for data in the volume.

- **Logical Space Reporting**

Whether the volume has logical space reporting configured. The value can be Enabled, Disabled, or Not applicable.

Logical space indicates the real size of the data that is being stored on the volume without applying the savings from using ONTAP storage efficiency technologies.

- **Logical Space Used %**

The percentage of logical space used by data in the volume based on the total available data capacity.

- **Logical Space Used**

The logical space used by data in the volume.

- **Move Status**

The current status of the volume move operation. The status can be In Progress, Paused, Failed, or Completed.

- **Type**

The volume type. The volume type can be Read-write or Data-protection, Load-sharing, or Data-cache.

- **Thin Provisioned**

Whether space guarantee is set for the selected volume. Valid values are Yes and No.

- **Deduplication**

Whether deduplication is enabled on the volume. The column displays either Enabled or Disabled.

- **Compression**

Whether compression is enabled on the volume. The column displays either Enabled or Disabled.

- **In Transition**

Whether the volume has completed transition or not.

- **SnapLock Type**

The SnapLock Type of the aggregate that contains the volume. The available options are Compliance, Enterprise, Non-SnapLock.

- **Local Snapshot Policy**

The local Snapshot copy policies for the volumes listed. The default policy name is Default.

- **Tiering Policy**

The tiering policy set on the volume. The policy takes affect only when the volume is deployed on a FabricPool aggregate:

- None - The data for this volume always remains on the performance tier.
- Snapshot-Only - Only Snapshot data is moved automatically to the cloud tier. All other data remains on the performance tier.
- Backup - On data protection volumes, all transferred user data starts in the cloud tier, but later client reads can cause hot data to move to the performance tier.
- Auto - Data on this volume is moved between the performance tier and the cloud tier automatically when ONTAP determines that the data is “hot” or “cold”.
- All - The data for this volume always remains on the cloud tier.

- **Caching Policy**

The caching policy that is associated with the selected volume. The policy provides information about how the Flash Pool caching occurs for the volume.

Cache policy	Description
Auto	Read caches all the metadata blocks and randomly read user data blocks, and write caches all the randomly overwritten user data blocks.
None	Does not cache any user data or metadata blocks.
All	Read caches all the user data blocks that are read and written. The policy does not perform any write caching.
All-Random Write	<p>This policy is a combination of the All and No Read-Random Write policies and performs the following actions:</p> <ul style="list-style-type: none"> • Read caches all the user data blocks that are read and written. • Write caches all the randomly overwritten user data blocks.
All Read	Read caches all the metadata, randomly read, and sequentially read user data blocks.
All Read-Random Write	<p>This policy is a combination of the All Read and No Read-Random Write policies and performs the following actions:</p> <ul style="list-style-type: none"> • Read caches all the metadata, randomly read, and sequentially read user data blocks. • Write caches all the randomly overwritten user data blocks.
All Read Random Write	Read caches all the metadata, randomly read, sequentially read, and randomly written user data blocks.
All Read Random Write-Random Write	<p>This policy is a combination of the All Read Random Write and No Read-Random Write policies and does the following:</p> <ul style="list-style-type: none"> • Read caches all the metadata, randomly read, and sequentially read, and randomly written user data blocks. • Write caches all the randomly overwritten user data blocks.
Meta	Read caches only metadata blocks.

Cache policy	Description
Meta-Random Write	This policy is a combination of the Meta and No Read-Random Write and does the following: Read caches only
No Read-Random Write	Write caches all the randomly overwritten user data blocks. The policy does not perform any read caching.
Random Read	Read caches all the metadata blocks and randomly read user data blocks.
Random Read-Write	Read caches all the metadata, randomly read, and randomly written user data blocks.
Random Read-Write-Random Write	This policy is a combination of the Random Read Write and No Read-Random Write policies and does the following: <ul style="list-style-type: none"> • Read caches all the metadata, randomly read, and randomly overwritten user data blocks. • Write caches all the randomly overwritten user data blocks.

• Cache Retention Priority

The cache retention priority for the volume. A cache retention priority defines how long the blocks of a volume will be in cache state in a Flash Pool once they become cold.

- Low
Cache the cold volume blocks for the lowest time
- Normal
Cache the cold volume blocks for the default time
- High
Cache the cold volume blocks for the highest time

• Encryption Type

The type of encryption that is applied to a volume.

- 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.

- **Aggregate**

The name of the aggregate on which the volume resides, or the number of aggregates on which the FlexGroup volume resides.

You can click the name to display details in the Aggregate details page. For FlexGroup volumes, you can click the number to display the aggregates that are used in the FlexGroup in the Aggregates page.

- **Node**

The name of the node to which the volume belongs, or the number of nodes on which the FlexGroup volume resides. You can view more details about the cluster node by clicking the node name.

You can click the node name to display details in the Node details page. For FlexGroup volumes, you can click the number to display the nodes that are used in the FlexGroup in the Nodes page.

- **Cluster**

The cluster that contains the destination volume. You can view more details about the cluster by clicking the cluster name.

- **Cluster FQDN**

The fully qualified domain name (FQDN) of the cluster.

- **Storage Class**

The storage class name. This column is displayed for Infinite Volume only.


- **Constituent Role**

The role name of the constituent. The roles can be Namespace, Data, or Namespace Mirror. This column is displayed for Infinite Volumes only.

Capacity: All Volumes view

The Capacity: All Volumes view enables you to view information about the capacity and utilization of volumes in a cluster. This information enables you to understand possible capacity risks and to view the configured, used, and unused capacity of volumes. Also, the information helps you to make decisions about enabling space-saving features such as deduplication and thin provisioning.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

If you want to analyze the latency and throughput of a specific object, click the more icon , then **Analyze Workload** and you can view performance and capacity charts on the Workload Analysis page.

See [Volume capacity fields](#) for descriptions of all the fields on this page.

Volume capacity fields

The following fields are available in the Capacity: All Volumes view and can be used in custom views and in reports.

- **Volume**

The volume name.

- **Daily Growth Rate %**

The growth rate that occurs every 24 hours in the volume.

- **Days To Full**

The estimated number of days remaining before the volume reaches full capacity.

- **Available Data %**

The available data capacity in a volume as a percentage.

- **Available Data Capacity**

The available data capacity in a volume.

- **Used Data %**

The used data in a volume as a percentage.

- **Used Data Capacity**

The used data capacity in a volume.

- **Total Data Capacity**

The total data capacity (used plus available) in a volume.

- **Logical Space Used**

The logical space used by data in this volume without applying the savings from using ONTAP storage efficiency technologies.

- **Snapshot Overflow %**

The percentage of the data space that is consumed by Snapshot copies.

- **Snapshot Reserve Available %**

The amount of space available for Snapshot copies in the volume as a percentage.

- **Snapshot Reserve Available Capacity**

The amount of space available for Snapshot copies in the volume.

- **Snapshot Reserve Used %**

The amount of space used by Snapshot copies in the volume as a percentage.

- **Snapshot Reserve Used Capacity**

The amount of space used by Snapshot copies in the volume.

- **Snapshot Reserve Total Capacity**

Displays the total Snapshot copy capacity in the volume.

- **Quota Committed Capacity**

The space reserved for quotas in the volume.

- **Quota Overcommitted Capacity**

The amount of space that can be used for quotas before the system generates the Volume Quota Overcommitted event.

- **Total Number Of Inodes**

The number of inodes in the volume.

- **Inode Utilization %**

The percentage of inode space used in the volume.

- **Thin Provisioned**

Whether space guarantee is set for the selected volume. Valid values are Yes and No.

- **Space Guarantee**

The storage guarantee option that is associated with the volume.

- **Autogrow**

Whether the volume automatically grows in size when it is out of space.

- **Snapshot Autodelete**

Whether automatic deletion of Snapshot copies is enabled or disabled.

- **Deduplication**

Whether deduplication is enabled or disabled for the volume.

- **Deduplication Space Savings**

The amount of space saved in a volume by using deduplication.

- **Compression**

Whether compression is enabled or disabled for the volume.

- **Compression Space Savings**

The amount of space saved in a volume by using compression.

- **State**

The state of the volume that is being exported.

- **Protection Role**

The protection role that is set for the volume.

- **SnapLock Type**

Whether the volume is a SnapLock or non-SnapLock volume.

- **SnapLock Expiry Date**

The SnapLock expiration date.

- **Tiering Policy**

The tiering policy set for the volume. Valid when deployed on FabricPool-enabled aggregates only.

- **Caching Policy**

The caching policy that is associated with the selected volume.

The policy provides information about how Flash Pool caching occurs for the volume. See the Health: All Volumes view for more information on caching policies.

- **Cache Retention Priority**

The priority used for retaining cached pools.

- **Storage VM**

The name of the storage virtual machine (SVM) that contains the volume.

- **Cluster**

The name of the cluster on which the volume resides. You can click the cluster name to navigate to that cluster's health details page.

- **Cluster FQDN**

The fully qualified domain name (FQDN) of the cluster.

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.



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.

Efficiency tab

The Efficiency tab displays information about the space saved in the volumes by using storage efficiency features such as deduplication, compression, and FlexClone volumes.

- **Deduplication**

- Enabled

Specifies whether deduplication is enabled or disabled on a volume.

- Space Savings

Displays the amount of space saved (in percentage, or in KB, MB, GB, and so on) in a volume by using deduplication.

- Last Run

Displays the time that has elapsed since the deduplication operation was last performed. Also specifies whether the deduplication operation was successful.

If the time elapsed exceeds a week, the timestamp representing when the operation was performed is displayed.

- 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.

- Status

Displays the current status of the deduplication operation. The status can be Idle, Initializing, Active, Undoing, Pending, Downgrading, or Disabled.

- 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.

- **Compression**

- Enabled

Specifies whether compression is enabled or disabled on a volume.

- Space Savings

Displays the amount of space saved (in percentage, or in KB, MB, GB, and so on) in a volume by using compression.

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 Virtual Machine

Displays the name of the storage virtual machine (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**

- **Deduplication**

Specifies whether deduplication is enabled or disabled for the selected volume.

- **Compression**

Specifies whether compression is enabled or disabled for the selected volume.

- **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 SnapMirror and SnapVault relationships 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.

- **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 transferred 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.

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.

Capacity: All Qtrees view

The Capacity: All Qtrees view enables you to view information about the capacity and utilization of qtrees in all clusters. This information enables you to understand possible capacity risks and also to view the configured and used disk percentage and number of files.

By default, objects in the view pages are sorted based on event criticality. Objects with critical events are listed first, and objects with warning events are listed second. This provides an immediate visual indication of issues

that must be addressed.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

You can customize capacity threshold settings for one or more qtrees by using the **Edit Threshold** button.

See [Qtree capacity fields](#) for descriptions of all the fields on this page.

Qtree capacity fields

The following fields are available in the Capacity: All Qtrees view and can be used in custom views and in reports.

- **Status**

Displays the current status of the qtree. The status can be Critical (❌), Error (⚠️), Warning (⚠️), or Normal (✅).

- **Qtree**

Displays the name of the qtree.

- **Volume**

Displays the name of the volume that contains the qtree.

You can click the volume name to view more information about the volume.

- **Quota Type**

If a quota is set for the qtree, specifies if the quota is for a User, User Group, or a Tree.

- **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.

- **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.

- **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.

- **SVM**

Displays the storage virtual machine (SVM) name containing the qtree.

- **Cluster**

Displays the name of the cluster containing the qtree.

- **Cluster FQDN**


Displays the fully qualified domain name (FQDN) of the cluster.

Health: All NFS Shares view

The Health: All NFS Shares view 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.

By default, objects on this page are sorted based on status. Objects with Errors are listed first, and objects that have Normal status are listed next. This provides an immediate visual indication of issues that must be addressed.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

If you want to analyze the latency and throughput of a specific object, click the more icon  , then **Analyze Workload** and you can view performance and capacity charts on the Workload Analysis page.

See [NFS Shares health fields](#) for descriptions of all the fields on this page.

NFS Shares health fields

The following fields are available in the Health: All NFS Shares view and can be used in custom views and in reports.

- **Status**

Displays the current status of the NFS shares. The status can be Error () or Normal () .

- **Mount 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.

- **Mount Path Active**

Displays whether the path to access the mounted volume is active or inactive.

- **Qtree**

Displays the name of the qtree to which the NFS export policy is applied.

- **Volume**

Displays the name of the volume to which the NFS export policy is applied.

- **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.

When you generate a report for the Health: All NFS Shares view, all rules that belong to the export policy are exported to the CSV or PDF file.

- **Rule Index**

Displays the rules associated with the export policy such as the authentication protocols and the access permission.

- **Access Protocols**

Displays the protocols that are enabled for the export policy rules.

- **Client Match**

Displays the clients that have permission to access data on the volumes.

- **Read Only Access**

Displays the authentication protocol used to read data on the volumes.

- **Read Write Access**

Displays the authentication protocol used to read or write data on the volumes.

- **Storage VM**

Displays the name of the SVM with NFS share policies.

- **Cluster**

Displays the name of the cluster.

- **Cluster FQDN**


Displays the fully qualified domain name (FQDN) of the cluster.

Health: All SMB Shares view

The Health: All SMB Shares view displays information about SMB/CIFS shares such as its status, the share name, junction path, containing objects, security settings, and export policies defined for the share.

By default, objects on this page are sorted based on status. Objects with Errors are listed first, and objects that have Normal status are listed next. This provides an immediate visual indication of issues that must be addressed.

The controls along the top of the page enable you to select a particular view (for health, performance, capacity, and so on), perform searches to locate specific objects, create and apply filters to narrow the list of displayed data, add/remove/reorder columns on the page, and export the data on the page to a `.csv` or `.pdf` document. After you have customized the page, you can save the results as a custom view and then schedule a report of this data to be generated and emailed on a regular basis.

If you want to analyze the latency and throughput of a specific object, click the more icon , then **Analyze Workload** and you can view performance and capacity charts on the Workload Analysis page.

See [SMB/CIFS Shares health fields](#) for descriptions of all the fields on this page.

SMB/CIFS Shares health fields

The following fields are available in the Health: All SMB Shares view and can be used in custom views and in reports.

- **View User Mapping button**

Launches the User Mapping dialog box.

You can view the details of user mapping for the SVM.

- **View ACL 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 ().

- **Name**

Displays the name of the CIFS share.

- **Path**

Displays the junction path on which the share is created.

- **Qtree**

Displays the name of the qtree to which the CIFS share is applied.

- **Volume**

Displays the name of the volume to which the CIFS share is applied.

- **Volume State**

Displays the state of the volume that is being exported. The state can be Offline, Online, Restricted, Mixed, or Unknown.

- 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.

- **Properties**

List the optional properties set when the share was created.

- **User**

The users who can access the share.

- **Permission**

The permissions users have on the share.

- **Security Style**

Displays the access permission for the volumes that are shared. 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 SVM, the value is displayed as Not Enabled.

- **Mount Path Active**

Displays whether the path to access the share is active or inactive.

- **NFS Equivalent**

Specifies whether there is an NFS equivalent for the share.

- **Storage VM**

Displays the name of the SVM with to which the CIFS share belongs.

- **Cluster**

Displays the name of the cluster.

- **Cluster FQDN**

Displays the fully qualified domain name (FQDN) of the cluster.

Export Policy Rules dialog box

The Export Policy Rules dialog box displays details about the export policies, authentication protocols, and access enabled on the volumes that belong to the storage virtual machine (SVM). You can use the filters to customize the display of information in the export policy rules list. By default, the information is sorted based on the index column.

- **Index**

Displays the index assigned to the export policy rules. It is a unique number.

- **Access Protocols**

Displays the protocols that are enabled for the export policy rules.

- **Client Match**

Displays the clients that have permission to access data on the volumes that belong to the SVM.

- **Read Only Access**

Displays the authentication protocol used to read data on the volumes that belong to the SVM.

- **Read Write Access**

Displays the authentication protocol used to read or write data on the volumes that belong to the SVM.

Snapshot Copies on a Volume dialog box

You can use the Snapshot Copies on a Volume dialog box to view the list of Snapshot copies. You can delete a Snapshot copy to conserve or free disk space, or if the copy is no longer required. You can also calculate the amount of disk space that can be reclaimed if one or more Snapshot copies are deleted.

List view

The list view displays, in tabular format, information about the Snapshot copies on the volume. You can use the column filters to customize the data that is displayed.

- **Snapshot Copy**

Displays the name of the Snapshot copy.

- **Used Space %**

Displays, in percentage, the total space used by the Snapshot copy in the volume.

- **Total Size**

Displays the total size of the Snapshot copy.

- **Created Time**

Displays the timestamp when the Snapshot copy was created.

- **Dependency**

Displays the applications that are dependent on the Snapshot copy. The possible values are SnapMirror, SnapVault, SnapLock, Dump, LUNs, Vclone, and Busy.

Command buttons

The command buttons enable you to perform the following tasks:

- **Calculate**

Enables you to calculate the space that can be reclaimed by deleting one or more Snapshot copies.

- **Delete Selected**

Deletes one or more Snapshot copies.

- **Close**

Closes the Snapshot Copies on a Volume dialog box.

- **Recalculate**

Enables you to calculate the space that can be reclaimed by deleting the selected Snapshot copies for FlexVol volumes. This button is not available for FlexGroup volumes.

The **Recalculate** button is enabled when you make any changes in the selection of the Snapshot copies.

Copyright information

Copyright © 2024 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.