



Tasks and information related to several workflows

Active IQ Unified Manager 9.7

NetApp
October 22, 2024

Table of Contents

- Tasks and information related to several workflows 1
 - Adding and reviewing notes about an event 1
 - Assigning events to specific users 1
 - Acknowledging and resolving events 2
 - Event details page 3
 - Description of event severity types 8
 - Description of event impact levels 9
 - Description of event impact areas 9
 - Cluster components and why they can be in contention 10
 - Adding alerts 12
 - Volume / Health details page 14
 - Storage VM / Health details page 30
 - Cluster / Health details page 44
 - Aggregate / Health details page 57
 - Adding users 67
 - Creating a database user 67
 - Definitions of user roles 68
 - Definitions of user types 69
 - Unified Manager user roles and capabilities 70
 - Supported Unified Manager CLI commands 71

Tasks and information related to several workflows

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

Adding and reviewing notes about an event

While addressing events, you can add information about how the issue is being addressed by using the Notes and Updates area in the Event details page. This information can enable another user who is assigned to address the event. You can also view information that was added by the user who last addressed an event, based on the recent timestamp.

Before you begin

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

Steps

1. In the left navigation pane, click **Events**.
2. From the **Event Management** inventory page, click the event for which you want to add the event-related information.
3. In the **Event** details page, add the required information in the **Notes and Updates** area.
4. Click **Post**.

Assigning events to specific users

You can assign unassigned events to yourself or to other users, including remote users. You can reassign assigned events to another user, if required. For example, when frequent issues occur on a storage object, you can assign the events for these issues to the user who manages that object.


Before you begin

- The user's name and email ID must be configured correctly.
- You must have the Operator, Application Administrator, or Storage Administrator role.

Steps

1. In the left navigation pane, click **Event Management**.
2. In the **Event Management** inventory page, select one or more events that you want to assign.

3. Assign the event by choosing one of the following options:

If you want to assign the event to...	Then do this...
Yourself	Click Assign To > Me .
Another user	<p>a. Click Assign To > Another user.</p> <p>b. In the Assign Owner dialog box, enter the user name, or select a user from the drop-down list.</p> <p>c. Click Assign.</p> <p>An email notification is sent to the user.</p> <div data-bbox="922 625 976 678"></div> <p data-bbox="1036 583 1430 716">If you do not enter a user name or select a user from the drop-down list, and click Assign, the event remains unassigned.</p>

Acknowledging and resolving events

You should acknowledge an event before you start working on the issue that generated the event so that you do not continue to receive repeat alert notifications. After you take corrective action for a particular event, you should mark the event as resolved.

Before you begin

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

About this task

You can acknowledge and resolve multiple events simultaneously.



You cannot acknowledge Information events.

Steps

1. In the left navigation pane, click **Event Management**.
2. From the events list, perform the following actions to acknowledge the events:

If you want to...	Do this...
Acknowledge and mark a single event as resolved	<ol style="list-style-type: none"> Click the event name. From the Event details page, determine the cause of the event. Click Acknowledge. Take appropriate corrective action. Click Mark As Resolved.
Acknowledge and mark multiple events as resolved	<ol style="list-style-type: none"> Determine the cause of the events from the respective Event details page. Select the events. Click Acknowledge. Take appropriate corrective actions. Click Mark As Resolved.

After the event is marked resolved, the event is moved to the resolved events list.

- In the **Notes and Updates** area, add a note about how you addressed the event, and then click **Post**.

Event details page

From the Event details page, you can view the details of a selected event, such as the event severity, impact level, impact area, and event source. You can also view additional information about possible remediations to resolve the issue.

- **Event Name**

The name of the event and the time the event was last seen.

For non-performance events, while the event is in the New or Acknowledged state the last seen information is not known and is therefore hidden.

- **Event Description**

A brief description of the event.

In some cases a reason for the event being triggered is provided in the event description.

- **Component in Contention**

For dynamic performance events, this section displays icons that represent the logical and physical components of the cluster. If a component is in contention, its icon is circled and highlighted red.

See [Cluster components and why they can be in contention](#) for a description of the components that are displayed here.

The Event Information, System Diagnosis, and Suggested Actions sections are described in other topics.

Command buttons

The command buttons enable you to perform the following tasks:

- **Notes icon**

Enables you to add or update a note about the event, and review all notes left by other users.

Actions menu

- **Assign to Me**

Assigns the event to you.

- **Assign to Others**

Opens the Assign Owner dialog box, which enables you to assign or reassign the event to other users.

When you assign an event to a user, the user's name and the time when the event was assigned are added in the events list for the selected events.

You can also unassign events by leaving the ownership field blank.

- **Acknowledge**

Acknowledges the selected events so that you do not continue to receive repeat alert notifications.

When you acknowledge an event, your user name and the time that you acknowledged the event are added in the events list (Acknowledged By) for the selected events. When you acknowledge an event, you take responsibility for managing that event.

- **Mark As Resolved**

Enables you to change the event state to Resolved.

When you resolve an event, your user name and the time that you resolved the event are added in the events list (Resolved By) for the selected events. After you have taken corrective action for the event, you must mark the event as resolved.

- **Add Alert**

Displays the Add Alert dialog box, which enables you to add an alert for the selected event.

What the Event Information section displays

You use the Event Information section on the Event details page to view the details about a selected event, such as the event severity, impact level, impact area, and event source.

Fields that are not applicable to the event type are hidden. You can view the following event details:

- **Event Trigger Time**

The time at which the event was generated.

- **State**

The event state: New, Acknowledged, Resolved, or Obsolete.

- **Obsoleted Cause**

The actions that caused the event to be obsoleted, for example, the issue was fixed.

- **Event Duration**

For active (new and acknowledged) events, this is the time between detection and the time when the event was last analyzed. For obsolete events, this is the time between detection and when the event was resolved.

This field is displayed for all performance events, and for other event types only after they have been resolved or obsoleted.

- **Last Seen**

The date and time at which the event was last seen as active.

For performance events this value may be more recent than the Event Trigger Time as this field is updated after each new collection of performance data as long as the event is active. For other types of events, when in the New or Acknowledged state, this content is not updated and the field is therefore hidden.

- **Severity**

The event severity: Critical (❌), Error (⚠️), Warning (⚠️), and Information (ℹ️).

- **Impact Level**

The event impact level: Incident, Risk, Event, or Upgrade.

- **Impact Area**

The event impact area: Availability, Capacity, Performance, Protection, Configuration, or Security.

- **Source**

The name of the object on which the event has occurred.

When viewing the details for a shared QoS policy event, up to three of the workload objects that are consuming the most IOPS or MBps are listed in this field.

You can click the source name link to display the health or performance details page for that object.

- **Source Annotations**

Displays the annotation name and value for the object to which the event is associated.

This field is displayed only for health events on clusters, SVMs, and volumes.

- **Source Groups**

Displays the names of all the groups of which the impacted object is a member.

This field is displayed only for health events on clusters, SVMs, and volumes.

- **Source Type**

The object type (for example, SVM, Volume, or Qtree) with which the event is associated.

- **On Cluster**

The name of the cluster on which the event occurred.

You can click the cluster name link to display the health or performance details page for that cluster.

- **Affected Objects Count**

The number of objects affected by the event.

You can click the object link to display the inventory page populated with the objects that are currently affected by this event.

This field is displayed only for performance events.

- **Affected Volumes**

The number of volumes that are being affected by this event.

This field is displayed only for performance events on nodes or aggregates.

- **Triggered Policy**

The name of the threshold policy that issued the event.

You can hover your cursor over the policy name to see the details of the threshold policy. For adaptive QoS policies the defined policy, block size, and allocation type (allocated space or used space) is also displayed.

This field is displayed only for performance events.

- **Rule Id**

For Active IQ platform events, this is the number of the rule that was triggered to generate the event.

- **Acknowledged by**

The name of the person who acknowledged the event and the time that the event was acknowledged.

- **Resolved by**

The name of the person who resolved the event and the time that the event was resolved.

- **Assigned to**

The name of the person who is assigned to work on the event.

- **Alert Settings**

The following information about alerts is displayed:

- If there are no alerts associated with the selected event, an **Add alert** link is displayed.

You can open the Add Alert dialog box by clicking the link.

- If there is one alert associated with the selected event, the alert name is displayed.

You can open the Edit Alert dialog box by clicking the link.

- If there is more than one alert associated with the selected event, the number of alerts is displayed.

You can open the Alert Setup page by clicking the link to view more details about these alerts.

Alerts that are disabled are not displayed.

- **Last Notification Sent**

The date and time at which the most recent alert notification was sent.

- **Send by**

The mechanism that was used to send the alert notification: email or SNMP trap.

- **Previous Script Run**

The name of the script that was executed when the alert was generated.

What the System Diagnosis section displays

The System Diagnosis section of the Event details page provides information that can help you diagnose issues that may have been responsible for the event.

This area is displayed only for some events.

Some performance events provide charts that are relevant to the particular event that has been triggered. Typically this includes an IOPS or MBps chart and a latency chart for the previous ten days. When arranged this way you can see which storage components are most affecting latency, or being affected by latency, when the event is active.

For dynamic performance events, the following charts are displayed:

- **Workload Latency** - Displays the history of latency for the top victim, bully, or shark workloads at the component in contention.
- **Workload Activity** - Displays details about the workload usage of the cluster component in contention.
- **Resource Activity** - Display historical performance statistics for the cluster component in contention.

Other charts are displayed when some cluster components are in contention.

Other events provide a brief description of the type of analysis the system is performing on the storage object. In some cases there will be one or more lines; one for each component that has been analyzed, for system-defined performance policies that analyze multiple performance counters. In this scenario, a green or red icon displays next to the diagnosis to indicate whether an issue was found, or not, in that particular diagnosis.

What the Suggested Actions section displays

The Suggested Actions section of the Event details page provides possible reasons for the event and suggests a few actions so that you can try to resolve the event on your own. The suggested actions are customized based on the type of event or type of threshold that has been breached.

This area is displayed only for some types of events.

In some cases there are **Help** links provided on the page that reference additional information for many suggested actions, including instructions for performing a specific action. Some of the actions may involve using Unified Manager, ONTAP System Manager, OnCommand Workflow Automation, ONTAP CLI commands, or a combination of these tools.

You should consider the actions suggested here as only a guidance in resolving this event. The action you take to resolve this event should be based on the context of your environment.

If you want to analyze the object and event in more detail, click the **Analyze Workload** button to display the Workload Analysis page.

There are certain events that Unified Manager can diagnose thoroughly and provide a single resolution. When available, those resolutions are displayed with a **Fix It** button. Click this button to have Unified Manager fix the issue causing the event.

For Active IQ platform events, this section may contain a link to a NetApp Knowledgebase article, when available, that describes the issue and possible resolutions. In sites with no external network access, a PDF of the Knowledgebase article is opened locally; the PDF is part of the rules file that you manually download to the Unified Manager instance.

Description of event severity types

Each event is associated with a severity type to help you prioritize the events that require immediate corrective action.

- **Critical**

A problem occurred that might lead to service disruption if corrective action is not taken immediately.

Performance critical events are sent from user-defined thresholds only.

- **Error**

The event source is still performing; however, corrective action is required to avoid service disruption.

- **Warning**

The event source experienced an occurrence that you should be aware of, or a performance counter for a cluster object is out of normal range and should be monitored to make sure it does not reach the critical severity. Events of this severity do not cause service disruption, and immediate corrective action might not be required.

Performance warning events are sent from user-defined, system-defined, or dynamic thresholds.

- **Information**

The event occurs when a new object is discovered, or when a user action is performed. For example, when any storage object is deleted or when there are any configuration changes, the event with severity type Information is generated.

Information events are sent directly from ONTAP when it detects a configuration change.

Description of event impact levels

Each event is associated with an impact level (Incident, Risk, Event, or Upgrade) to help you prioritize the events that require immediate corrective action.

- **Incident**

An incident is a set of events that can cause a cluster to stop serving data to the client and run out of space for storing data. Events with an impact level of Incident are the most severe. Immediate corrective action should be taken to avoid service disruption.

- **Risk**

A risk is a set of events that can potentially cause a cluster to stop serving data to the client and run out of space for storing data. Events with an impact level of Risk can cause service disruption. Corrective action might be required.

- **Event**

An event is a state or status change of storage objects and their attributes. Events with an impact level of Event are informational and do not require corrective action.

- **Upgrade**

Upgrade events are a specific type of event reported from the Active IQ platform. These events identify issues where the resolution requires you to upgrade ONTAP software, node firmware, or operating system software (for security advisories). You may want to perform immediate corrective action for some of these issues, whereas other issues may be able to wait until your next scheduled maintenance.

Description of event impact areas

Events are categorized into six impact areas (availability, capacity, configuration, performance, protection, and security) to enable you to concentrate on the types of events for which you are responsible.

- **Availability**

Availability events notify you if a storage object goes offline, if a protocol service goes down, if an issue with storage failover occurs, or if an issue with hardware occurs.

- **Capacity**

Capacity events notify you if your aggregates, volumes, LUNs, or namespaces are approaching or have reached a size threshold, or if the rate of growth is unusual for your environment.

- **Configuration**

Configuration events inform you of the discovery, deletion, addition, removal, or renaming of your storage objects. Configuration events have an impact level of Event and a severity type of Information.

- **Performance**

Performance events notify you of resource, configuration, or activity conditions on your cluster that might adversely affect the speed of data storage input or retrieval on your monitored storage objects.

- **Protection**

Protection events notify you of incidents or risks involving SnapMirror relationships, issues with destination capacity, problems with SnapVault relationships, or issues with protection jobs. Any ONTAP object (especially aggregates, volumes, and SVMs) that host secondary volumes and protection relationships are categorized in the protection impact area.

- **Security**

Security events notify you of how secure your ONTAP clusters, storage virtual machines (SVMs), and volumes are based on parameters defined in the [NetApp Security Hardening Guide for ONTAP 9](#).

Additionally, this area includes upgrade events that are reported from the Active IQ platform.

Cluster components and why they can be in contention

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

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

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

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

- **Network**

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

- **Network Processing**

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

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

- **QoS Limit Max**

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

- **QoS Limit Min**

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

- **Cluster Interconnect**

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

- **Data Processing**

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

- **Volume Activation**

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

- **MetroCluster Resources**

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

- **Aggregate or SSD Aggregate Ops**

Represents the storage aggregate on which the workloads are running. If the aggregate component is in contention, it means high utilization on the aggregate is impacting the latency of one or more workloads. An

aggregate consists of all HDDs, or a mix of HDDs and SSDs (a Flash Pool aggregate). An “SSD Aggregate” consists of all SSDs (an all-flash aggregate), or a mix of SSDs and a cloud tier (a FabricPool aggregate).

- **Cloud Latency**

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

- **Sync SnapMirror**

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

Adding alerts

You can configure alerts to notify you when a particular event is generated. You can configure alerts for a single resource, for a group of resources, or for events of a particular severity type. You can specify the frequency with which you want to be notified and associate a script to the alert.

Before you begin

- You must have configured notification settings such as the user email address, SMTP server, and SNMP trap host to enable the Active IQ Unified Manager server to use these settings to send notifications to users when an event is generated.
- You must know the resources and events for which you want to trigger the alert, and the user names or email addresses of the users that you want to notify.
- If you want to have a script execute based on the event, you must have added the script to Unified Manager by using the Scripts page.
- You must have the Application Administrator or Storage Administrator role.

About this task

You can create an alert directly from the Event details page after receiving an event in addition to creating an alert from the Alert Setup page, as described here.

Steps

1. In the left navigation pane, click **Storage Management > Alert Setup**.
2. In the **Alert Setup** page, click **Add**.
3. In the **Add Alert** dialog box, click **Name**, and enter a name and description for the alert.
4. Click **Resources**, and select the resources to be included in or excluded from the alert.

You can set a filter by specifying a text string in the **Name contains** field to select a group of resources.

Based on the text string that you specify, the list of available resources displays only those resources that match the filter rule. The text string that you specify is case-sensitive.

If a resource conforms to both the include and exclude rules that you have specified, the exclude rule takes precedence over the include rule, and the alert is not generated for events related to the excluded resource.

5. Click **Events**, and select the events based on the event name or event severity type for which you want to trigger an alert.



To select more than one event, press the Ctrl key while you make your selections.

6. Click **Actions**, and select the users that you want to notify, choose the notification frequency, choose whether an SNMP trap will be sent to the trap receiver, and assign a script to be executed when an alert is generated.



If you modify the email address that is specified for the user and reopen the alert for editing, the Name field appears blank because the modified email address is no longer mapped to the user that was previously selected. Also, if you modified the email address of the selected user from the Users page, the modified email address is not updated for the selected user.

You can also choose to notify users through SNMP traps.

7. Click **Save**.

Example of adding an alert

This example shows how to create an alert that meets the following requirements:

- Alert name: HealthTest
- Resources: includes all volumes whose name contains “abc” and excludes all volumes whose name contains “xyz”
- Events: includes all critical health events
- Actions: includes “sample@domain.com”, a “Test” script, and the user has to be notified every 15 minutes

Perform the following steps in the Add Alert dialog box:

1. Click **Name**, and enter `HealthTest` in the **Alert Name** field.
2. Click **Resources**, and in the Include tab, select **Volumes** from the drop-down list.
 - a. Enter `abc` in the **Name contains** field to display the volumes whose name contains “abc”.
 - b. Select **All Volumes whose name contains 'abc'** from the Available Resources area, and move it to the Selected Resources area.
 - c. Click **Exclude**, and enter `xyz` in the **Name contains** field, and then click **Add**.
3. Click **Events**, and select **Critical** from the Event Severity field.
4. Select **All Critical Events** from the Matching Events area, and move it to the Selected Events area.
5. Click **Actions**, and enter `sample@domain.com` in the Alert these users field.
6. Select **Remind every 15 minutes** to notify the user every 15 minutes.

You can configure an alert to repeatedly send notifications to the recipients for a specified time. You should determine the time from which the event notification is active for the alert.

7. In the Select Script to Execute menu, select **Test** script.

8. Click **Save**.

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.

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.

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.

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.

Adding users

You can add local users or database users by using the Users page. You can also add remote users or groups that belong to an authentication server. You can assign roles to these users and, based on the privileges of the roles, users can manage the storage objects and data with Unified Manager, or view the data in a database.

Before you begin

- You must have the Application Administrator role.
- To add a remote user or group, you must have enabled remote authentication and configured your authentication server.
- If you plan to configure SAML authentication so that an identity provider (IdP) authenticates users accessing the graphical interface, make sure these users are defined as “remote” users.

Access to the UI is not allowed for users of type “local” or “maintenance” when SAML authentication is enabled.

About this task

If you add a group from Windows Active Directory, then all direct members and nested subgroups can authenticate to Unified Manager, unless nested subgroups are disabled. If you add a group from OpenLDAP or other authentication services, then only the direct members of that group can authenticate to Unified Manager.

Steps

1. In the left navigation pane, click **General > Users**.
2. On the **Users** page, click **Add**.
3. In the **Add User** dialog box, select the type of user that you want to add, and enter the required information.

When entering the required user information, you must specify an email address that is unique to that user. You must avoid specifying email addresses that are shared by multiple users.

4. Click **Add**.

Creating a database user

To support a connection between Workflow Automation and Unified Manager, or to access database views, you must first create a database user with the Integration Schema or Report Schema role in the Unified Manager web UI.

Before you begin

You must have the Application Administrator role.

About this task

Database users provide integration with Workflow Automation and access to report-specific database views. Database users do not have access to the Unified Manager web UI or the maintenance console, and cannot execute API calls.

Steps

1. In the left navigation pane, click **General > Users**.
2. In the **Users** page, click **Add**.
3. In the **Add User** dialog box, select **Database User** in the **Type** drop-down list.
4. Type a name and password for the database user.
5. In the **Role** drop-down list, select the appropriate role.

If you are...	Choose this role
Connecting Unified Manager with Workflow Automation	Integration Schema
Accessing reporting and other database views	Report Schema

6. Click **Add**.

Definitions of user roles

The maintenance user or Application Administrator assigns a role to every user. Each role contains certain privileges. The scope of activities that you can perform in Unified Manager depends on the role you are assigned and which privileges the role contains.

Unified Manager includes the following predefined user roles:

- **Operator**

Views storage system information and other data collected by Unified Manager, including histories and capacity trends. This role enables the storage operator to view, assign, acknowledge, resolve, and add notes for the events.

- **Storage Administrator**

Configures storage management operations within Unified Manager. This role enables the storage administrator to configure thresholds and to create alerts and other storage management-specific options and policies.

- **Application Administrator**

Configures settings unrelated to storage management. This role enables the management of users, security certificates, database access, and administrative options, including authentication, SMTP, networking, and AutoSupport.



When Unified Manager is installed on Linux systems, the initial user with the Application Administrator role is automatically named “umadmin”.

- **Integration Schema**

This role enables read-only access to Unified Manager database views for integrating Unified Manager with OnCommand Workflow Automation (WFA).

- **Report Schema**

This role enables read-only access to reporting and other database views directly from the Unified Manager database. The databases that can be viewed include:

- netapp_model_view
- netapp_performance
- ocum
- ocum_report
- ocum_report_birt
- opm
- scalemonitor

Definitions of user types

A user type specifies the kind of account the user holds and includes remote users, remote groups, local users, database users, and maintenance users. Each of these types has its own role, which is assigned by a user with the role of Administrator.

Unified Manager user types are as follows:

- **Maintenance user**

Created during the initial configuration of Unified Manager. The maintenance user then creates additional users and assigns roles. The maintenance user is also the only user with access to the maintenance console. When Unified Manager is installed on a Red Hat Enterprise Linux or CentOS system, the maintenance user is given the user name “umadmin.”

- **Local user**

Accesses the Unified Manager UI and performs functions based on the role given by the maintenance user or a user with the Application Administrator role.

- **Remote group**

A group of users that access the Unified Manager UI using the credentials stored on the authentication server. The name of this account should match the name of a group stored on the authentication server. All users within the remote group are given access to the Unified Manager UI using their individual user credentials. Remote groups can perform functions according to their assigned roles.

- **Remote user**

Accesses the Unified Manager UI using the credentials stored on the authentication server. A remote user

performs functions based on the role given by the maintenance user or a user with the Application Administrator role.

- **Database user**

Has read-only access to data in the Unified Manager database, has no access to the Unified Manager web interface or the maintenance console, and cannot execute API calls.

Unified Manager user roles and capabilities

Based on your assigned user role, you can determine which operations you can perform in Unified Manager.

The following table displays the functions that each user role can perform:

Function	Operator	Storage Administrator	Application Administrator	Integration Schema	Report Schema
View storage system information	•	•	•	•	•
View other data, such as histories and capacity trends	•	•	•	•	•
View, assign, and resolve events	•	•	•		
View storage service objects, such as SVM associations and resource pools	•	•	•		
View threshold policies	•	•	•		
Manage storage service objects, such as SVM associations and resource pools		•	•		
Define alerts		•	•		

Function	Operator	Storage Administrator	Application Administrator	Integration Schema	Report Schema
Manage storage management options		•	•		
Manage storage management policies		•	•		
Manage users			•		
Manage administrative options			•		
Define threshold policies			•		
Manage database access			•		
Manage integration with WFA and provide access to the database views				•	
Provide read-only access to database views					•
Schedule and save reports		•	•		

Supported Unified Manager CLI commands

As a storage administrator you can use the CLI commands to perform queries on the storage objects; for example, on clusters, aggregates, volumes, qtrees, and LUNs. You can use the CLI commands to query the Unified Manager internal database and the ONTAP database. You can also use CLI commands in scripts that are executed at the beginning or end of an operation or are executed when an alert is triggered.

All commands must be preceded with the command `um cli login` and a valid user name and password for authentication.

CLI command	Description	Output
um cli login -u <username> [-p <password>]	Logs in to the CLI. Because of security implications, you should enter only the user name following the “-u” option. When used in this manner you will be prompted for the password, and the password will not be captured in the history or process table. The session expires after three hours from the time of login, after which the user must login again.	Displays the corresponding message.
um cli logout	Logs out of the CLI.	Displays the corresponding message.
um help	Displays all first level subcommands.	Displays all first level subcommands.
um run cmd [-t <timeout> >] <cluster> <command>	The simplest way to run a command on one or more hosts. Mainly used for alert scripting to get or perform an operation on ONTAP. The optional timeout argument sets a maximum time limit (in seconds) for the command to complete on the client. The default is 0 (wait forever).	As received from ONTAP.
um run query <sql command>	Executes an SQL query. Only queries that read from the database are allowed. Any update, insert, or delete operations are not supported.	Results are displayed in a tabular form. If an empty set is returned, or if there is any syntax error or bad request, it displays the appropriate error message.

CLI command	Description	Output
<pre>um datasource add -u <username> -P <password> [-t <protocol>] [-p <port>] <hostname-or-ip></pre>	<p>Adds a datasource to the list of managed storage systems. A datasource describes how connections to storage systems are made. The options -u (username) and -P (password) must be specified when adding a datasource. The option -t (protocol) specifies the protocol used to communicate with the cluster (http or https). If the protocol is not specified, then both protocols will be attempted. The option -p (port) specifies the port used to communicate with the cluster. If the port is not specified, then the default value of the appropriate protocol will be attempted. This command can be executed only by the storage admin.</p>	<p>Prompts for the user accept the certificate and prints the corresponding message.</p>
<pre>um datasource list [<datasource-id>]</pre>	<p>Displays the datasources for managed storage systems.</p>	<p>Displays the following values in tabular format: ID Address Port, Protocol Acquisition Status, Analysis Status, Communication status, Acquisition Message, and Analysis Message.</p>
<pre>um datasource modify [-h <hostname-or-ip>] [-u <username>] [-P <password>] [-t <protocol>] [-p <port>] <datasource-id></pre>	<p>Modifies one or more datasource options. Can be executed only by the storage admin.</p>	<p>Displays the corresponding message.</p>
<pre>um datasource remove <datasource-id></pre>	<p>Removes the datasource (cluster) from Unified Manager.</p>	<p>Displays the corresponding message.</p>
<pre>um option list [<option> ..]</pre>	<p>Lists options.</p>	<p>Displays the following values in tabular format: Name, Value, Default Value, and Requires Restart.</p>
<pre>um option set <option- name>=<option-value> [<option-name>=<option- value> ...]</pre>	<p>Sets one or more options. The command can be executed only by the storage admin.</p>	<p>Displays the corresponding message.</p>

CLI command	Description	Output
<code>um version</code>	Displays the Unified Manager software version.	Version ("9.6")
<code>um lun list [-q] [-ObjectType <object-id>]</code>	<p>Lists the LUNs after filtering on the specified object. -q is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, or svm. For example: <code>um lun list -cluster 1</code></p> <p>In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the LUNs within the cluster with ID 1.</p>	Displays the following values in tabular format: ID and LUN path.
<code>um svm list [-q] [-ObjectType <object-id>]</code>	<p>Lists the SVMs after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, or svm. For example: <code>um svm list -cluster 1</code></p> <p>In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the SVMs within the cluster with ID 1.</p>	Displays the following values in tabular format: Name and Cluster ID.
<code>um qtree list [-q] [-ObjectType <object-id>]</code>	<p>Lists the qtrees after filtering on the specified object. -q is applicable for all commands to show no header. ObjectType can be lun, qtree, cluster, volume, quota, or svm. For example: <code>um qtree list -cluster 1</code></p> <p>In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the qtrees within the cluster with ID 1.</p>	Displays the following values in tabular format: Qtree ID and Qtree Name.

CLI command	Description	Output
<pre>um disk list [-q] [- ObjectType <object-id>]</pre>	<p>Lists the disks after filtering on the specified object. ObjectType can be disk, aggr, node, or cluster. For example: <code>um disk list -cluster 1</code></p> <p>In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the disks within the cluster with ID 1.</p>	<p>Displays the following values in tabular format ObjectType and object-id.</p>
<pre>um cluster list [-q] [- ObjectType <object-id>]</pre>	<p>Lists the clusters after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, lun, qtree, volume, quota, or svm. For example: <code>um cluster list -aggr 1</code></p> <p>In this example, "-aggr" is the objectType and "1" is the objectId. The command lists the cluster to which the aggregate with ID 1 belongs.</p>	<p>Displays the following values in tabular format: Name, Full Name, Serial Number, Datasource Id, Last Refresh Time, and Resource Key.</p>
<pre>um cluster node list [-q] [- ObjectType <object-id>]</pre>	<p>Lists the cluster nodes after filtering on the specified object. ObjectType can be disk, aggr, node, or cluster. For example: <code>um cluster node list -cluster 1</code></p> <p>In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the nodes within the cluster with ID 1.</p>	<p>Displays the following values in tabular format Name and Cluster ID.</p>
<pre>um volume list [-q] [- ObjectType <object-id>]</pre>	<p>Lists the volumes after filtering on the specified object. ObjectType can be lun, qtree, cluster, volume, quota, svm, or aggregate. For example: <code>um volume list -cluster 1</code></p> <p>In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the volumes within the cluster with ID 1.</p>	<p>Displays the following values in tabular format Volume ID and Volume Name.</p>

CLI command	Description	Output
um quota user list [-q] [-ObjectType <object-id>]	<p>Lists the quota users after filtering on the specified object. ObjectType can be qtree, cluster, volume, quota, or svm. For example: <code>um quota user list -cluster 1</code></p> <p>In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the quota users within the cluster with ID 1.</p>	Displays the following values in tabular format ID, Name, SID and Email.
um aggr list [-q] [-ObjectType <object-id>]	<p>Lists the aggregates after filtering on the specified object. ObjectType can be disk, aggr, node, cluster, or volume. For example: <code>um aggr list -cluster 1</code></p> <p>In this example, "-cluster" is the objectType and "1" is the objectId. The command lists all the aggregates within the cluster with ID 1.</p>	Displays the following values in tabular format Aggr ID, and Aggr Name.
um event ack <event-ids>	Acknowledges one or more events.	Displays the corresponding message.
um event resolve <event-ids>	Resolves one or more events.	Displays the corresponding message.
um event assign -u <username> <event-id>	Assigns an event to a user.	Displays the corresponding message.
um event list [-s <source>] [-S <event-state-filter-list>..] [<event-id> ..]	Lists the events generated by the system or user. Filters events based on source, state, and IDs.	Displays the following values in tabular format Source, Source type, Name, Severity, State, User and Timestamp.
um backup restore -f <backup_file_path_and_name>	Restores a database backup using .7z files.	Displays the corresponding message.

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.