



Common Unified Manager health workflows and tasks

Active IQ Unified Manager 9.7

NetApp
October 22, 2024

Table of Contents

- Common Unified Manager health workflows and tasks 1
 - Monitoring and troubleshooting data availability 2
 - Resolving capacity issues 8
 - Managing health thresholds 10
 - Managing cluster security objectives 15
 - Configuring backup and restore operations 25
 - Managing scripts 33
 - Managing and monitoring groups 36
 - Prioritizing storage object events using annotations 45
 - What a Unified Manager maintenance window is 53
 - Managing SAML authentication settings 56
 - Sending a Unified Manager support bundle to technical support 62
 - Tasks and information related to several workflows 67

Common Unified Manager health workflows and tasks

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

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

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

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

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

- Diagnosing and managing availability issues

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

- Configuring and monitoring performance incidents

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

- Diagnosing and managing volume capacity issues

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

- Configuring, monitoring, and diagnosing protection relationship issues

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

- Creating backup files and restoring data from backup files.
- Associating storage objects with annotations

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

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

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

Monitoring and troubleshooting data availability

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

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

Scanning for and resolving storage failover interconnect link down conditions

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


Before you begin

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

About this task

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

Steps

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

All past events related to storage failover conditions are displayed.

In this scenario, the Unified Manager displays the event, “Storage Failover Interconnect One or More Links

Down” in its Availability Incidents section.

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

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

The Event details page for that event is displayed.

- b. On the **Event** details page, you can perform one or more of the following tasks:
 - Review the error message in the Cause field and evaluate the issue. [Performing corrective action for storage failover interconnect links down](#)
 - Assign the event to an administrator. [Assigning events](#)
 - Acknowledge the event. [Acknowledging and resolving events](#)

Performing corrective action for storage failover interconnect links down

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

Before you begin

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

About this task

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

```
Event: Storage Failover Interconnect One or More Links Down
```

```
Summary
```

```
Severity: Warning
```

```
State: New
```

```
Impact Level: Risk
```

```
Impact Area: Availability
```

```
Source: aardvark
```

```
Source Type: Node
```

```
Acknowledged By:
```

```
Resolved By:
```

```
Assigned To:
```

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

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

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

Steps

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

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

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

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

After you finish

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

Resolving volume offline issues

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

Before you begin

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

About this task

Volumes might be reported offline for several reasons:

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

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

Steps

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

The Event details page for the availability incident is displayed.

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

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

- Check the Notes and Updates area for any indication that the SVM administrator has deliberately taken the volume in question offline.
- Click the source of the event, in this case the volume that is reported offline, to get more information about that volume. [Performing corrective action for volume offline conditions](#)
- Assign the event to an administrator. [Assigning events](#)
- Acknowledge the event or, if appropriate, mark it as resolved. [Acknowledging and resolving events](#)

Performing diagnostic actions for volume offline conditions

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

Before you begin

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

About this task

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

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

Choices

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

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

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

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

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

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

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

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

Before you begin

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

About this task

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


Steps

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


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

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

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

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

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

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

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

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

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

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

After you finish

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

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

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

Before you begin

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


About this task

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

Steps

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

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

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

The information for that SVM is shown in a table.

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

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

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

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

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

The Nodes tab on the Cluster / Health details page displays the state of the HA pair of nodes to which the

SVM root volume's hosting node belongs. The state of the node is indicated in the display.

After you finish

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

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

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

Before you begin

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

About this task

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

Steps

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

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

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

The Cluster / Health details page is displayed.

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

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

After you finish

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

Resolving capacity issues

This workflow provides an example of how you can resolve a capacity issue. In this scenario, you are an administrator or operator and you access the Unified

ManagerDashboard page to see if any of the monitored storage objects have capacity issues. You want to determine the possible cause of and resolution to the problem.

Before you begin

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

About this task

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

Steps

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

The Event details page for the error is displayed.

2. From the **Event** details page, you can perform one or more of the following tasks:
 - Review the error message in the Cause field and click the suggestions under Suggested Remedial Actions to review descriptions of possible remediations. [Performing suggested remedial actions for a full volume](#)
 - Click the object name, in this case a volume, in the Source field to get details about the object. [Volume details page](#)
 - Look for notes that might have been added about this event. [Adding and reviewing notes associated with an event](#)
 - Add a note to the event. [Adding and reviewing notes associated with an event](#)
 - Assign the event to another user. [Assigning events](#)
 - Acknowledge the event. [Acknowledging and resolving events](#)
 - Mark the event as resolved. [Acknowledging and resolving events](#)

Performing suggested remedial actions for a full volume

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

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

About this task

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

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

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

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

Steps

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

You decide to resize the volume.

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

Managing health thresholds

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

What storage capacity health thresholds are

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

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

Configuring global health threshold settings

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

About this task

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

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

Choices

- [Configuring global aggregate health threshold values](#)

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

- [Configuring global volume health threshold values](#)

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

- [Configuring global qtree health threshold values](#)

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

- [Editing lag health threshold settings for unmanaged protection relationships](#)

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

Configuring global aggregate health threshold values

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

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

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



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

Steps

1. In the left navigation pane, click **Event Thresholds > Aggregate**.

2. Configure the appropriate threshold values for capacity, growth, and Snapshot copies.
3. Click **Save**.

Configuring global volume health threshold values

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

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

Steps

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

Configuring global qtree health threshold values

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

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

Steps

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

Configuring lag threshold settings for unmanaged protection relationships

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

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

Steps

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

Editing individual aggregate health threshold settings

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

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

size.

Steps

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

Editing individual volume health threshold settings

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

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

Steps

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

Editing individual qtree health threshold settings

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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



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

Managing cluster security objectives

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

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

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

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

What security criteria is being evaluated

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

Some of the security checks include:

- whether a cluster is using a secure authentication method, such as SAML
- whether peered clusters have their communication encrypted
- whether a storage VM has its audit log enabled
- whether your volumes have software or hardware encryption enabled

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



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

Cluster compliance categories

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

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

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

Parameter	Description	Recommendation	Affects Cluster Compliance
Global FIPS	Indicates if Global FIPS (Federal Information Processing Standard) 140-2 compliance mode is enabled or disabled. When FIPS is enabled, TLSv1 and SSLv3 are disabled, and only TLSv1.1 and TLSv1.2 are allowed.	Enabled	Yes
Telnet	Indicates if Telnet access to the system is enabled or disabled. NetApp recommends Secure Shell (SSH) for secure remote access.	Disabled	Yes
Insecure SSH Settings	Indicates if SSH uses insecure ciphers, for example ciphers beginning with *cbc.	No	Yes
Login Banner	Indicates if the Login banner is enabled or disabled for users accessing the system.	Enabled	Yes

Parameter	Description	Recommendation	Affects Cluster Compliance
Cluster Peering	Indicates if communication between peered clusters is encrypted or unencrypted. Encryption must be configured on both the source and destination clusters for this parameter to be considered compliant.	Encrypted	Yes
Network Time Protocol	Indicates if the cluster has one or more configured NTP servers. For redundancy and best service NetApp recommends that you associate at least three NTP servers with the cluster.	Configured	Yes
OCSP	Indicates if there are applications in ONTAP that are not configured with OCSP (Online Certificate Status Protocol) and therefore communications are not encrypted. The non-compliant applications are listed.	Enabled	No
Remote Audit Logging	Indicates if log forwarding (Syslog) is encrypted or not encrypted.	Encrypted	Yes
AutoSupport HTTPS Transport	Indicates if HTTPS is used as the default transport protocol for sending AutoSupport messages to NetApp support.	Enabled	Yes

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

SVM compliance categories

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

Parameter	Description	Recommendation	Affects SVM Compliance
Audit Log	Indicates if Audit logging is enabled or disabled.	Enabled	Yes
Insecure SSH Settings	Indicates if SSH uses insecure ciphers, for example ciphers beginning with <code>cbc*</code> .	No	Yes
Login Banner	Indicates if the Login banner is enabled or disabled for users accessing SVMs on the system.	Enabled	Yes
LDAP Encryption	Indicates if LDAP Encryption is enabled or disabled.	Enabled	No
NTLM Authentication	Indicates if NTLM Authentication is enabled or disabled.	Enabled	No
LDAP Payload Signing	Indicates if LDAP Payload Signing is enabled or disabled.	Enabled	No
CHAP Settings	Indicates if CHAP is enabled or disabled.	Enabled	No
Kerberos V5	Indicates if Kerberos V5 authentication is enabled or disabled.	Enabled	No

Volume compliance categories

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




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

Parameter	Description
Software Encrypted	Displays the number of volumes that are protected using NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) software encryption solutions.
Hardware Encrypted	Displays the number of volumes that are protected using NetApp Storage Encryption (NSE) hardware encryption.
Software and Hardware Encrypted	Displays the number of volumes that are protected by both software and hardware encryption.
Not Encrypted	Displays the number of volumes that are not encrypted.

What does not compliant mean

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

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

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

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

Viewing high-level cluster security status

The Security panel on the Unified ManagerDashboard shows high-level security status for all clusters or for a single cluster, depending on your current view.

Steps

1. In the left navigation pane, click **Dashboard**.
2. Depending on whether you want to view security status for all monitored clusters or for a single cluster, select **All Clusters** or select a single cluster from the drop-down menu.
3. View the **Security** panel to see the overall status.

This panel displays:

- a list of the security events received in the past 24 hours

- a link from each of these events to the Event details page
 - a link so that you can view all active security events in the Event Management inventory page
 - the cluster security status (number of clusters that are compliant or not compliant)
 - the SVM security status (number of SVMs that are compliant or not compliant)
 - the volume encryption status (number of volumes that are encrypted or not encrypted)
4. Click the right-arrow at the top of the panel to view security details in the **Security** page.

Viewing detailed security status for clusters and SVMs

The Security page shows high-level security status for all clusters, and detailed security status for individual clusters. The detailed cluster status includes cluster compliance, SVM compliance, and volume encryption compliance.

Steps

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

The Security page displays the following information:

- the cluster security status (number of clusters that are compliant or not compliant)
 - the SVM security status (number of SVMs that are compliant or not compliant)
 - the volume encryption status (number of volumes that are encrypted or not encrypted)
 - the cluster authentication methods being used on each cluster
4. Refer to the [NetApp Security Hardening Guide for ONTAP 9](#) for instructions on how to make all of your clusters, SVMs, and volumes compliant with NetApp security recommendations.

Viewing security events that may require software or firmware updates

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

Before you begin

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

About this task

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

Steps

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

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

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

The page displays all active upgrade security events.

Viewing how user authentication is being managed on all clusters

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

Steps

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

Results

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

Viewing the encryption status of all volumes

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

Before you begin

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

About this task

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

- Software - Volumes that are protected using NetApp Volume Encryption (NVE) or NetApp Aggregate Encryption (NAE) software encryption solutions.
- Hardware - Volumes that are protected using NetApp Storage Encryption (NSE) hardware encryption.

- Software and Hardware - Volumes that are protected by both software and hardware encryption.
- None - Volumes that are not encrypted.

Steps

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

Viewing all active security events

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

Before you begin

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

Steps

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

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

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

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

Adding alerts for security events

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

Disabling specific security events

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

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

Security events

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

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

Configuring backup and restore operations


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

What a database backup is

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

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

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



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

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

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

Note that you can restore a backup only on the same version of Unified Manager. For example, if you created a backup on Unified Manager 9.4, the backup can be restored only on Unified Manager 9.4 systems.

Configuring database backup settings

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

Before you begin

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

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

- When Unified Manager is installed on a Linux system, verify that the “jboss” user has write permissions to the backup directory.
- You should not schedule backup operations to occur immediately after a new cluster has been added while Unified Manager is collecting 15 days of historical performance data.

About this task

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

Steps

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

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

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

What a database restore is

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

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

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

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

- Virtual appliance to virtual appliance
- Virtual appliance to Red Hat Enterprise Linux or CentOS
- Red Hat Enterprise Linux to Red Hat Enterprise Linux or CentOS
- Windows to Windows

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



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

Virtual appliance backup and restore process overview

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

Because the Unified Manager backup operation on the virtual appliance does not provide a way to move the backup file from the vApp, the following tasks enable you to complete a backup of the virtual appliance:

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

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

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

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

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

Restoring a database backup on a virtual machine

In case of data loss or data corruption, you can use the restore feature to restore Unified Manager to the previous stable state with minimal loss. You can restore the Unified Manager database on a virtual machine by using the Unified Manager maintenance console.

Before you begin

- You must have the maintenance user credentials.
- The Unified Manager backup files must be on the local system.
- The backup files must be of .7z type.

About this task

Backup compatibility is platform and version dependent. You can restore a backup from a virtual appliance to another virtual appliance, or from a virtual appliance to a Red Hat Enterprise Linux or CentOS system.



When performing a restore operation on a different virtual appliance than the system from which the original backup file was created, the maintenance user name and password on the new vApp must be the same as the credentials from the original vApp.

Steps

1. In the vSphere client, locate the Unified Manager virtual machine, and then select the **Console** tab.
2. Click in the console window, and then log in to the maintenance console using your user name and password.

3. In the **Main Menu**, enter the number for the **System Configuration** option.
4. In the **System Configuration Menu**, enter the number for the **Restore from a Unified Manager Backup** option.
5. When prompted, enter the absolute path of the backup file.

```
Bundle to restore from: opt/netapp/data/ocum-  
backup/UM_9.4.N151112.0947_backup_unix_02-25-2018-11-41.7z
```

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

After you finish

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

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

Restoring a database backup on a Linux system

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

Before you begin

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

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

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

About this task

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



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

Steps

1. If you are performing a restore onto a new server, after installing Unified Manager do not launch the UI or configure any clusters, users, or authentication settings when the installation is complete. The backup file populates this information during the restore process.
2. Using Secure Shell, connect to the IP address or fully qualified domain name of the Unified Manager system.
3. Log in to the system with the maintenance user (umadmin) name and password.
4. Enter the command `maintenance_console` and press Enter.
5. In the maintenance console **Main Menu**, enter the number for the **System Configuration** option.
6. In the **System Configuration Menu**, enter the number for the **Restore from a Unified Manager Backup** option.
7. When prompted, enter the absolute path of the backup file.

```
Bundle to restore from: /data/ocum-  
backup/UM_9.4.N151113.1348_backup_rhel_02-20-2018-04-45.7z
```

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

After you finish

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

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

Restoring a database backup on Windows

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

Before you begin

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

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

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

About this task

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



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

Steps

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

```
Bundle to restore from:  
\\ProgramData\NetApp\OnCommandAppData\ocum\backup\UM_9.4.N151118.2300_bac  
kup_windows_02-20-2018-02-51.7z
```

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

After you finish

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

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

Migrating a Unified Manager virtual appliance to a Linux system

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


Before you begin

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

About this task

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

Steps

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

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

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

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

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

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

8. At the command prompt, restore the backup:

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

```
um backup restore -f /UM_9.4.N151113.1348_backup_unix_02-12-2018-04-16.7z
```
9. After the restore operation completes, log in to the Unified Manager web UI.

After you finish

You should perform the following tasks:

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

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

Managing scripts

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

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

How scripts work with alerts

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

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

The script uses the following arguments for execution:

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

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

Example for obtaining arguments from scripts

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

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

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

Adding scripts

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

Before you begin

- You must have created and saved the scripts that you want to add to the Unified Manager server.
- The supported file formats for scripts are Perl, Shell, PowerShell, and .bat files.
 - For Perl scripts, Perl must be installed on the Unified Manager server. If Perl was installed after Unified Manager, you must restart the Unified Manager server.
 - For PowerShell scripts, the appropriate PowerShell execution policy must be set on the server so that the scripts can be executed.



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

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

About this task

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



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

Steps

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

Deleting scripts

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

Before you begin

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

Steps

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

Testing script execution

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

Before you begin

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

Steps

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

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

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

The script is executed with the “-test” argument, and a notification alert is sent to the email addresses that

were specified when the alert was created.

Managing and monitoring groups

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

Understanding groups

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

What a group is

A group is a dynamic collection of heterogeneous storage objects (clusters, SVMs, or volumes). You can create groups in Unified Manager to easily manage a set of storage objects. The members in a group might change, depending on the storage objects that are monitored by Unified Manager at a point in time.

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

How group rules work for groups

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

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

Conditions

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

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

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

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

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

- Operand

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

- Operator

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

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

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

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

- Value

The value field changes based on the operand selected.

Example of a group rule with conditions

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

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

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

Condition groups

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

Example of a group rule with condition groups

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

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

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

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

How group actions work on storage objects

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

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

Example of unique group actions

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

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

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

Adding groups

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

Editing groups

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

Deleting groups

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

Before you begin

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

Steps

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

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

Adding group rules

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

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

7. Click **Add**.

Example for creating a group rule

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

- Specify a name for the group rule.
- Select the object type as storage virtual machine (SVM).
- Select a group from the list of groups.
- In the Conditions section, select **Object Name** as the operand.
- Select **Contains** as the operator.
- Enter the value as `svm_data`.
- Click **Add condition group**.
- Select **Object Name** as the operand.
- Select **Contains** as the operator.
- Enter the value as `vol`.
- Click **Add condition**.
- Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **critical** as the value in step 10.
- Click **Add** to create the condition for the group rule.

Editing group rules

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

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



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

4. Click **Save**.

Deleting group rules

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

Adding group actions

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

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

4. From the **Group** menu, select a group for which you want to configure the action.
5. From the **Action Type** menu, select an action type.

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

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

Editing group actions

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

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

Configuring volume health thresholds for groups

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

7. Enter the required values for the health threshold.
8. Click **Add**.

Deleting group actions

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

Reordering group actions

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

1. In the left navigation pane, click **Storage Management > Groups**.
2. In the **Group Actions** tab, click **Reorder**.
3. In the **Reorder Group Actions** dialog box, drag and drop the rows to rearrange the sequence of group actions as required.
4. Click **Save**.

Prioritizing storage object events using annotations

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

Understanding more about annotations

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

What annotations are

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

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

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

How annotation rules work in Unified Manager

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

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

Conditions

You can create multiple condition groups, and each condition group can have one or more conditions. You can apply all the defined condition groups in an annotation rule of an annotation in order to annotate storage objects.

Conditions within a condition group are executed using logical AND. All the conditions in a condition group

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

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

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

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

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

Example of an annotation rule with conditions

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

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

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

Condition groups

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

Example of an annotation rule with condition groups

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

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

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

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

Description of predefined annotation values

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

- **Data-priority:Mission critical**

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

- **Data-priority:High**

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

- **Data-priority:Low**

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

Adding annotations dynamically

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

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

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

Adding values to annotations

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

You cannot add values to predefined annotations.

Steps

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

The value that you specify must be unique for the selected annotation.

4. Click **Add**.

Deleting annotations

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

Before you begin

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

Steps

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

The details of the selected annotation are displayed.

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

Viewing the annotation list and details

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

Steps

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

Deleting values from annotations

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

Before you begin

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

About this task

You cannot delete values from predefined annotations.

Steps

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

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

Creating annotation rules

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotation Rules** tab, click **Add**.
3. In the **Add Annotation Rule** dialog box, specify a name for the annotation rule.
4. In the **Target Object Type** field, select the type of storage object that you want to annotate.
5. In the **Apply Annotation** fields, select the annotation and annotation value that you want to use.
6. In the **Conditions** section, perform the appropriate action to create a condition, a condition group, or both:

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

7. Click **Add**.

Example of creating an annotation rule

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

1. Specify a name for the annotation rule.
2. Select the target object type as storage virtual machine (SVM).
3. Select an annotation from the list of annotations, and specify a value.
4. In the Conditions section, select **Object Name** as the operand.
5. Select **Contains** as the operator.
6. Enter the value as `svm_data`.
7. Click **Add condition group**.
8. Select **Object Name** as the operand.
9. Select **Contains** as the operator.
10. Enter the value as `vol`.
11. Click **Add condition**.

12. Repeat steps 8 through 10 by selecting **data-priority** as the operand in step 8, **Is** as the operator in step 9, and **mission-critical** as the value in step 10.
13. Click **Add**.

Adding annotations manually to individual storage objects

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

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

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

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

Editing annotation rules

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

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

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

4. Click **Save**.

Configuring conditions for annotation rules

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

Steps

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

Example of configuring a condition for an annotation rule

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

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

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

Deleting annotation rules

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

1. In the left navigation pane, click **Storage Management > Annotations**.
2. In the **Annotation Rules** tab, select the annotation rule that you want to delete, and then click **Delete**.
3. In the **Warning** dialog box, click **Yes** to confirm the deletion.

Reordering annotation rules

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

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

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

Steps

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

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

What a Unified Manager maintenance window is

You define a Unified Manager maintenance window to suppress events and alerts for a specific timeframe when you have scheduled cluster maintenance and you do not want to receive a flood of unwanted notifications.

When the maintenance window starts, an “Object Maintenance Window Started” event is posted to the Event Management inventory page. This event is obsoleted automatically when the maintenance window ends.

During a maintenance window the events related to all objects on that cluster are still generated, but they do not appear in any of the UI pages, and no alerts or other types of notification are sent for these events. You can, however, view the events that were generated for all storage objects during a maintenance window by selecting one of the View options on the Event Management inventory page.

You can schedule a maintenance window to be initiated in the future, you can change the start and end times for a scheduled maintenance window, and you can cancel a scheduled maintenance window.

Scheduling a maintenance window to disable cluster event notifications

If you have a planned downtime for a cluster, for example, to upgrade the cluster or to move one of the nodes, you can suppress the events and alerts that would normally be generated during that timeframe by scheduling a Unified Manager maintenance window.

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

During a maintenance window, the events related to all objects on that cluster are still generated, but they do not appear in the event page, and no alerts or other types of notification are sent for these events.

The time you enter for the maintenance window is based on the time at the Unified Manager server.

Steps

1. In the left navigation pane, click **Storage Management > Cluster Setup**.
2. In the **Maintenance Mode** column for the cluster, select the slider button and move it to the right.

The calendar window is displayed.

3. Select the start and end date and time for the maintenance window and click **Apply**.

The message “Scheduled” appears next to the slider button.

Results

When the start time is reached the cluster goes into maintenance mode and an “Object Maintenance Window Started” event is generated.

Changing or canceling a scheduled maintenance window

If you have configured a Unified Manager maintenance window to occur in the future, you can change the start and end times or cancel the maintenance window from occurring.

Before you begin

You must have the Application Administrator or Storage Administrator role.

About this task

Canceling a currently running maintenance window is useful if you have completed cluster maintenance before the scheduled maintenance window end time and you want to start receiving events and alerts from the cluster again.

Steps

1. In the left navigation pane, click **Storage Management > Cluster Setup**.
2. In the **Maintenance Mode** column for the cluster:

If you want to...	Perform this step...
Change the timeframe for a scheduled maintenance window	<ol style="list-style-type: none">a. Click the text "Scheduled" next to the slider button.b. Change the start and/or end date and time and click Apply.
Extend the length of an active maintenance window	<ol style="list-style-type: none">a. Click the text "Active" next to the slider button.b. Change the end date and time and click Apply.
Cancel a scheduled maintenance window	Select the slider button and move it to the left.
Cancel an active maintenance window	Select the slider button and move it to the left.

Viewing events that occurred during a maintenance window

If necessary, you can view the events that were generated for all storage objects during a Unified Manager maintenance window. Most events will appear in the Obsolete state once the maintenance window has completed and all system resources are back up and running.

Before you begin

At least one maintenance window must have completed before any events are available.

About this task

Events that occurred during a maintenance window do not appear on the Event Management inventory page by default.

Steps

1. In the left navigation pane, click **Events**.

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

2. From the **View** pane, select the option **All events generated during maintenance**.

The list of events triggered during the last 7 days from all maintenance window sessions and from all clusters are displayed.

3. If there have been multiple maintenance windows for a single cluster, you can click the **Triggered Time** calendar icon and select the period of time for the maintenance window events that you are interested in viewing.

Managing SAML authentication settings

After you have configured remote authentication settings, you can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Note that only remote users will have access to the Unified Manager graphical user interface after SAML authentication has been enabled. Local users and Maintenance users will not be able to access the UI. This configuration does not impact users who access the maintenance console.

Identity provider requirements

When configuring Unified Manager to use an identity provider (IdP) to perform SAML authentication for all remote users, you need to be aware of some required configuration settings so that the connection to Unified Manager is successful.

You must enter the Unified Manager URI and metadata into the IdP server. You can copy this information from the Unified Manager SAML Authentication page. Unified Manager is considered the service provider (SP) in the Security Assertion Markup Language (SAML) standard.

Supported encryption standards

- Advanced Encryption Standard (AES): AES-128 and AES-256
- Secure Hash Algorithm (SHA): SHA-1 and SHA-256

Validated identity providers

- Shibboleth
- Active Directory Federation Services (ADFS)

ADFS configuration requirements

- You must define three claim rules in the following order that are required for Unified Manager to parse ADFS SAML responses for this relying party trust entry.

Claim rule	Value
SAM-account-name	Name ID
SAM-account-name	urn:oid:0.9.2342.19200300.100.1.1

Claim rule	Value
Token groups — Unqualified Name	urn:oid:1.3.6.1.4.1.5923.1.5.1.1

- You must set the authentication method to “Forms Authentication” or users may receive an error when logging out of Unified Manager . Follow these steps:
 - a. Open the ADFS Management Console.
 - b. Click on the Authentication Policies folder on the left tree view.
 - c. Under Actions on the right, click Edit Global Primary Authentication Policy.
 - d. Set the Intranet Authentication Method to “Forms Authentication” instead of the default “Windows Authentication”.
- In some cases login through the IdP is rejected when the Unified Manager security certificate is CA-signed. There are two workarounds to resolve this issue:
 - Follow the instructions identified in the link to disable the revocation check on the ADFS server for chained CA cert associated relying party:

<http://www.torivar.com/2016/03/22/adfs-3-0-disable-revocation-check-windows-2012-r2/>
 - Have the CA server reside within the ADFS server to sign the Unified Manager server cert request.

Other configuration requirements

- The Unified Manager clock skew is set to 5 minutes, so the time difference between the IdP server and the Unified Manager server cannot be more than 5 minutes or authentication will fail.

Enabling SAML authentication

You can enable Security Assertion Markup Language (SAML) authentication so that remote users are authenticated by a secure identity provider (IdP) before they can access the Unified Manager web UI.

Before you begin

- You must have configured remote authentication and verified that it is successful.
- You must have created at least one Remote User, or a Remote Group, with the Application Administrator role.
- The Identity provider (IdP) must be supported by Unified Manager and it must be configured.
- You must have the IdP URL and metadata.
- You must have access to the IdP server.

About this task

After you have enabled SAML authentication from Unified Manager, users cannot access the graphical user interface until the IdP has been configured with the Unified Manager server host information. So you must be prepared to complete both parts of the connection before starting the configuration process. The IdP can be configured before or after configuring Unified Manager.

Only remote users will have access to the Unified Manager graphical user interface after SAML authentication

has been enabled. Local users and Maintenance users will not be able to access the UI. This configuration does not impact users who access the maintenance console, the Unified Manager commands, or ZAPIs.



Unified Manager is restarted automatically after you complete the SAML configuration on this page.

Steps

1. In the left navigation pane, click **General > SAML Authentication**.
2. Select the **Enable SAML authentication** checkbox.

The fields required to configure the IdP connection are displayed.

3. Enter the IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP server.

If the IdP server is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URI to populate the IdP Metadata field automatically.

4. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.

You can configure the IdP server with this information at this time.

5. Click **Save**.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

6. Click **Confirm and Logout** and Unified Manager is restarted.

Results

The next time authorized remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the IdP login page instead of the Unified Manager login page.

After you finish

If not already completed, access your IdP and enter the Unified Manager server URI and metadata to complete the configuration.



When using ADFS as your identity provider, the Unified Manager GUI does not honor the ADFS timeout and will continue to work until the Unified Manager session timeout is reached. When Unified Manager is deployed on Windows, Red Hat, or CentOS, you can change the GUI session timeout using the following Unified Manager CLI command: `um option set absolute.session.timeout=00:15:00` This command sets the Unified Manager GUI session timeout to 15 minutes.

Changing the identity provider used for SAML authentication

You can change the identity provider (IdP) that Unified Manager uses to authenticate remote users.

Before you begin

- You must have the IdP URL and metadata.
- You must have access to the IdP.

About this task

The new IdP can be configured before or after configuring Unified Manager.

Steps

1. In the left navigation pane, click **General > SAML Authentication**.
2. Enter the new IdP URI and the IdP metadata required to connect the Unified Manager server to the IdP.

If the IdP is accessible directly from the Unified Manager server, you can click the **Fetch IdP Metadata** button after entering the IdP URL to populate the IdP Metadata field automatically.

3. Copy the Unified Manager metadata URI, or save the metadata to an XML text file.
4. Click **Save Configuration**.

A message box displays to confirm that you want to change the configuration.

5. Click **OK**.

After you finish

Access the new IdP and enter the Unified Manager server URI and metadata to complete the configuration.

The next time authorized remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the new IdP login page instead of the old IdP login page.

Updating SAML authentication settings after Unified Manager security certificate change

Any change to the HTTPS security certificate installed on the Unified Manager server requires that you update the SAML authentication configuration settings. The certificate is updated if you rename the host system, assign a new IP address for the host system, or manually change the security certificate for the system.

About this task

After the security certificate is changed and the Unified Manager server is restarted, SAML authentication will not function and users will not be able to access the Unified Manager graphical interface. You must update the SAML authentication settings on both the IdP server and on the Unified Manager server to re-enable access to the user interface.

Steps

1. Log into the maintenance console.
2. In the **Main Menu**, enter the number for the **Disable SAML authentication** option.

A message displays to confirm that you want to disable SAML authentication and restart Unified Manager.

3. Launch the Unified Manager user interface using the updated FQDN or IP address, accept the updated server certificate into your browser, and log in using the maintenance user credentials.
4. In the **Setup/Authentication** page, select the **SAML Authentication** tab and configure the IdP connection.
5. Copy the Unified Manager host metadata URI, or save the host metadata to an XML text file.
6. Click **Save**.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

7. Click **Confirm and Logout** and Unified Manager is restarted.
8. Access your IdP server and enter the Unified Manager server URI and metadata to complete the configuration.

Identity provider	Configuration steps
ADFS	<ol style="list-style-type: none"> a. Delete the existing relying party trust entry in the ADFS management GUI. b. Add a new relying party trust entry using the <code>saml_sp_metadata.xml</code> from the updated Unified Manager server. c. Define the three claim rules that are required for Unified Manager to parse ADFS SAML responses for this relying party trust entry. d. Restart the ADFS Windows service.
Shibboleth	<ol style="list-style-type: none"> a. Update the new FQDN of Unified Manager server into the <code>attribute-filter.xml</code> and <code>relying-party.xml</code> files. b. Restart the Apache Tomcat web server and wait for port 8005 to come online.

9. Log in to Unified Manager and verify that SAML authentication works as expected through your IdP.

Disabling SAML authentication

You can disable SAML authentication when you want to stop authenticating remote users through a secure identity provider (IdP) before they can log into the Unified Manager web UI. When SAML authentication is disabled, the configured directory service providers, such as Active Directory or LDAP, perform sign-on authentication.

About this task

After you disable SAML authentication, Local users and Maintenance users will be able to access the graphical user interface in addition to configured Remote users.

You can also disable SAML authentication using the Unified Manager maintenance console if you do not have access to the graphical user interface.



Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

1. In the left navigation pane, click **General > SAML Authentication**.
2. Uncheck the **Enable SAML authentication** checkbox.
3. Click **Save**.

A message box displays to confirm that you want to complete the configuration and restart Unified Manager.

4. Click **Confirm and Logout** and Unified Manager is restarted.

Results

The next time remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the Unified Manager login page instead of the IdP login page.

After you finish

Access your IdP and delete the Unified Manager server URI and metadata.

Disabling SAML authentication from the maintenance console

You may need to disable SAML authentication from the maintenance console when there is no access to the Unified Manager GUI. This could happen in cases of misconfiguration or if the IdP is not accessible.

Before you begin

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

About this task

When SAML authentication is disabled, the configured directory service providers, such as Active Directory or LDAP, perform sign-on authentication. Local users and Maintenance users will be able to access the graphical user interface in addition to configured Remote users.

You can also disable SAML authentication from the Setup/Authentication page in the UI.



Unified Manager is restarted automatically after SAML authentication is disabled.

Steps

1. Log into the maintenance console.
2. In the **Main Menu**, enter the number for the **Disable SAML authentication** option.

A message displays to confirm that you want to disable SAML authentication and restart Unified Manager.

3. Type **y**, and then press Enter and Unified Manager is restarted.

Results

The next time remote users attempt to access the Unified Manager graphical interface they will enter their credentials in the Unified Manager login page instead of the IdP login page.

After you finish

If required, access your IdP and delete the Unified Manager server URL and metadata.

Sending a Unified Manager support bundle to technical support

This workflow shows you how to generate, retrieve, and send a support bundle to technical support using the Unified Manager maintenance console. You should send a support bundle when the issue that you have requires more detailed diagnosis and troubleshooting than an AutoSupport message provides.

About this task

For more information about the maintenance console, see [Maintenance console menus](#).

Unified Manager stores two generated support bundles at one time.

Accessing the maintenance console

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

Before you begin

You must have installed and configured Unified Manager.

About this task

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



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

Steps

1. Follow these steps to access the maintenance console:

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

The Unified Manager maintenance console menu is displayed.

Generating a support bundle

You can generate a support bundle, containing full diagnostic information, so that you can then retrieve it and send it to technical support for troubleshooting help. Because some types of data can use a large amount of cluster resources or take a long time to complete, you can specify data types to include or exclude in the support bundle.

Before you begin

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

About this task

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



On Windows systems, the command `supportbundle.bat` is no longer supported to generate a support bundle.

Steps

- In the maintenance console **Main Menu**, select **Support/Diagnostics**.
- Select **Generate Support Bundle**.

3. Select or deselect the following data types to include or exclude in the support bundle:

- **database dump**

A dump of the MySQL Server database.

- **heap dump**

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

- **acquisition recordings**

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



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

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

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

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

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

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

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

For example, enter `20170101` for January 1, 2017. Enter `n` if you do not want performance statistics to be included.

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

You can enter a number from 1 through 10.

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

7. Select **Generate Support Bundle**.

The generated support bundle resides in the `/support` directory.

After you finish

After generating the support bundle, you can retrieve it using an SFTP client or by using UNIX or Linux CLI commands. On Windows installations you can use Remote Desktop (RDP) to retrieve the support bundle.

The generated support bundle resides in the `/support` directory on VMware systems, in `/opt/netapp/data/support/` on Linux systems, and in

Retrieving the support bundle using a Windows client

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

Before you begin

You must be the maintenance user to perform this task.

You must use a tool that supports SCP or SFTP.

Steps

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

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

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

Related information

Filezilla - <https://filezilla-project.org/>

WinSCP - <http://winscp.net>

Retrieving the support bundle using a UNIX or Linux client

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

Before you begin

You must be the maintenance user to perform this task.

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

Steps

1. Access the CLI through Telnet or the console, using your Linux client server.
2. Access the `/support` directory.

3. Retrieve the support bundle and copy it to the local directory using the following command:

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

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

4. Enter the maintenance user password.

Examples

The following example uses SCP to retrieve the support bundle:

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

The following example uses SFTP to retrieve the support bundle:

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

Sending a support bundle to technical support

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

Before you begin

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

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

Steps

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

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

Tasks and information related to several workflows

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

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	<ol style="list-style-type: none">a. Click Assign To > Another user.b. In the Assign Owner dialog box, enter the user name, or select a user from the drop-down list.c. Click Assign. <p>An email notification is sent to the user.</p> <div data-bbox="922 781 977 835"></div> <p data-bbox="1036 743 1429 877">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"> a. Click the event name. b. From the Event details page, determine the cause of the event. c. Click Acknowledge. d. Take appropriate corrective action. e. Click Mark As Resolved.
Acknowledge and mark multiple events as resolved	<ol style="list-style-type: none"> a. Determine the cause of the events from the respective Event details page. b. Select the events. c. Click Acknowledge. d. Take appropriate corrective actions. e. Click Mark As Resolved.

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

3. 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	•	•	•	•	•

Function	Operator	Storage Administrator	Application Administrator	Integration Schema	Report Schema
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		•	•		
Manage storage management options		•	•		
Manage storage management policies		•	•		
Manage users			•		
Manage administrative options			•		
Define threshold policies			•		
Manage database access			•		

Function	Operator	Storage Administrator	Application Administrator	Integration Schema	Report Schema
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
<code>um cli login -u <username> [-p <password>]</code>	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.
<code>um cli logout</code>	Logs out of the CLI.	Displays the corresponding message.
<code>um help</code>	Displays all first level subcommands.	Displays all first level subcommands.

CLI command	Description	Output
<code>um run cmd [-t <timeout>] <cluster> <command></code>	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.
<code>um run query <sql command></code>	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.
<code>um datasource add -u <username> -P <password> [-t <protocol>] [-p <port>] <hostname-or-ip></code>	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.	Prompts for the user accept the certificate and prints the corresponding message.
<code>um datasource list [<datasource-id>]</code>	Displays the datasources for managed storage systems.	Displays the following values in tabular format: ID Address Port, Protocol Acquisition Status, Analysis Status, Communication status, Acquisition Message, and Analysis Message.

CLI command	Description	Output
<code>um datasource modify [-h <hostname-or-ip>] [-u <username>] [-P <password>] [-t <protocol>] [-p <port>] <datasource-id></code>	Modifies one or more datasource options. Can be executed only by the storage admin.	Displays the corresponding message.
<code>um datasource remove <datasource-id></code>	Removes the datasource (cluster) from Unified Manager.	Displays the corresponding message.
<code>um option list [<option> ..]</code>	Lists options.	Displays the following values in tabular format: Name, Value, Default Value, and Requires Restart.
<code>um option set <option-name>=<option-value> [<option-name>=<option-value> ...]</code>	Sets one or more options. The command can be executed only by the storage admin.	Displays the corresponding message.
<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.

CLI command	Description	Output
<pre>um qtree list [-q] [-Objectype <object-id>]</pre>	<p>Lists the qtrees after filtering on the specified object. -q is applicable for all commands to show no header. Objectype 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 objectype and "1" is the objectid. The command lists all the qtrees within the cluster with ID 1.</p>	<p>Displays the following values in tabular format: Qtree ID and Qtree Name.</p>
<pre>um disk list [-q] [-Objectype <object-id>]</pre>	<p>Lists the disks after filtering on the specified object. Objectype can be disk, aggr, node, or cluster. For example: <code>um disk list -cluster 1</code></p> <p>In this example, "-cluster" is the objectype 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 Objectype and object-id.</p>
<pre>um cluster list [-q] [-Objectype <object-id>]</pre>	<p>Lists the clusters after filtering on the specified object. Objectype 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 objectype 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] [-Objectype <object-id>]</pre>	<p>Lists the cluster nodes after filtering on the specified object. Objectype 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 objectype 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>

CLI command	Description	Output
<code>um volume list [-q] [-ObjectType <object-id>]</code>	<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>	Displays the following values in tabular format Volume ID and Volume Name.
<code>um quota user list [-q] [-ObjectType <object-id>]</code>	<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.
<code>um aggr list [-q] [-ObjectType <object-id>]</code>	<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.
<code>um event ack <event-ids></code>	Acknowledges one or more events.	Displays the corresponding message.
<code>um event resolve <event-ids></code>	Resolves one or more events.	Displays the corresponding message.
<code>um event assign -u <username> <event-id></code>	Assigns an event to a user.	Displays the corresponding message.
<code>um event list [-s <source>] [-S <event-state-filter-list>..] [<event-id> ..]</code>	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.

CLI command	Description	Output
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.