



Monitor and manage storage

Active IQ Unified Manager 9.14

NetApp
September 13, 2024

This PDF was generated from https://docs.netapp.com/us-en/active-iq-unified-manager/storage-mgmt/concept_introduction_to_unified_manager_health_monitoring.html on September 13, 2024. Always check docs.netapp.com for the latest.

Table of Contents

- Monitor and manage storage 1
 - Introduction to Active IQ Unified Manager 1
 - Understanding the user interface 4
 - Monitoring and managing clusters from the dashboard 11
 - Managing clusters 21
 - Monitoring VMware virtual infrastructure 27
 - Provisioning and managing workloads. 36
 - Managing and monitoring MetroCluster configurations 51
 - Managing quotas 58
 - Troubleshooting 65

Monitor and manage storage

Introduction to Active IQ Unified Manager

Active IQ Unified Manager (formerly OnCommand Unified Manager) enables you to monitor and manage the health and performance of your ONTAP storage systems from a single interface.

Unified Manager provides the following features:

- Discovery, monitoring, and notifications for systems that are installed with ONTAP software.
- Dashboard to show capacity, security, and performance health of the environment.
- Enhanced alerts, events, and threshold infrastructure.
- Displays detailed graphs that plot workload activity over time; including IOPS (operations), MBps (throughput), latency (response time), utilization, performance capacity, and cache ratio.
- Identifies workloads that are overusing cluster components and the workloads whose performance is impacted by the increased activity.
- Provides suggested corrective actions that can be performed to address certain incidents and events, and a "Fix It" button for some events so you can resolve the issue immediately.
- Integrates with OnCommand Workflow Automation to execute automated protection workflows.
- Ability to create new workloads, such as a LUN or file share, directly from Unified Manager and assign a Performance Service Level to define the performance and storage objectives for the users accessing the application using that workload.

Introduction to Active IQ Unified Manager health monitoring

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

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

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

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

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

Physical and logical capacity

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

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

Capacity measurement units

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



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

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

Introduction to Active IQ Unified Manager performance monitoring

Active IQ Unified Manager (formerly OnCommand Unified Manager) provides performance monitoring capabilities and event root-cause analysis for systems that are running NetApp ONTAP software.

Unified Manager helps you to identify workloads that are overusing cluster components and decreasing the performance of other workloads on the cluster. By defining performance threshold policies you can also specify maximum values for certain performance counters so that events are generated when the threshold is breached. Unified Manager alerts you about these performance events so that you can take corrective action, and bring performance back to normal levels of operation. You can view and analyze events in the Unified Manager UI.

Unified Manager monitors the performance of two types of workloads:

- User-defined workloads

These workloads consist of FlexVol volumes and FlexGroup volumes that you have created in your cluster.

- System-defined workloads

These workloads consist of internal system activity.

Using Unified Manager REST APIs

Active IQ Unified Manager provides you with REST APIs to view the information about monitoring and managing your storage environment. APIs also allow provisioning and managing storage objects based on policies.

You can also execute ONTAP APIs on all ONTAP-managed clusters by using the API gateway supported by Unified Manager.

For information about Unified Manager REST APIs, see [Getting started with Active IQ Unified Manager REST APIs](#).

What the Unified Manager server does

The Unified Manager server infrastructure consists of a data collection unit, a database, and an application server. It provides infrastructure services such as discovery, monitoring, role-based access control (RBAC), auditing, and logging.

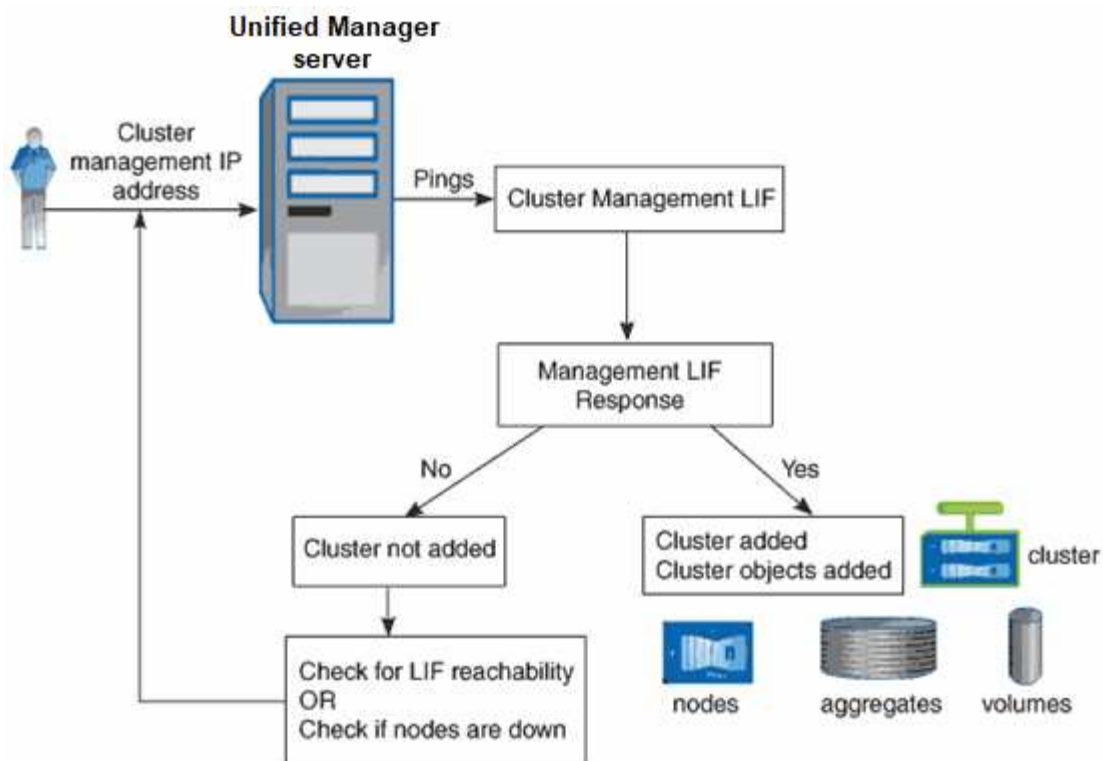
Unified Manager collects cluster information, stores the data in the database, and analyzes the data to see if there are any cluster issues.

How the discovery process works

After you have added the cluster to Unified Manager, the server discovers the cluster objects and adds them to its database. Understanding how the discovery process works helps you to manage your organization's clusters and their objects.

The default monitoring interval is 15 minutes: if you have added a cluster to Unified Manager server, it takes 15 minutes to display the cluster details in the Unified Manager UI.

The following image illustrates the discovery process in Active IQ Unified Manager:



Understanding the user interface

The Unified Manager user interface mainly consists of a dashboard that provides an at-a-glance view of the objects that are monitored. The user interface also provides access to viewing all the cluster objects.

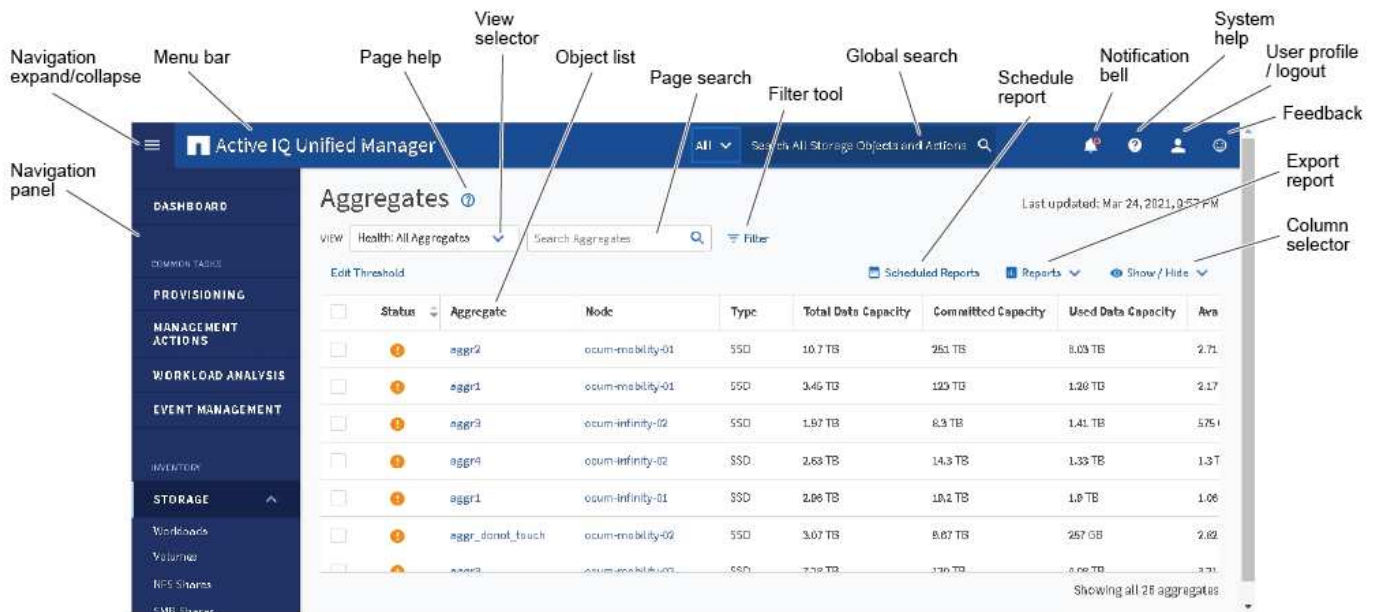
You can select a preferred view and use the action buttons as necessary. Your screen configuration is saved in a workspace so that all of the functionality you require is available when you start Unified Manager. However, when you navigate from one view to another, and then navigate back, the view might not be the same.

Typical window layouts

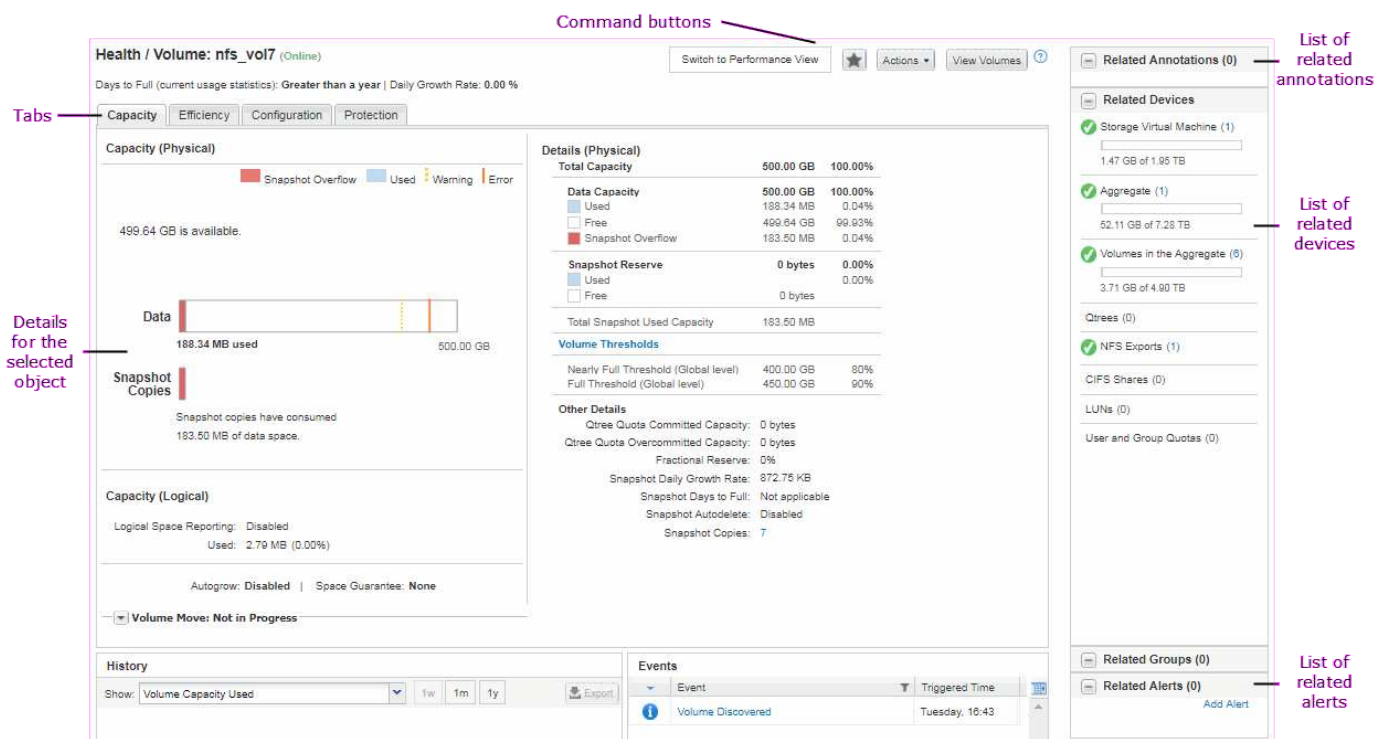
Understanding the typical window layouts helps you to navigate and use Active IQ Unified Manager effectively. Most Unified Manager windows are similar to one of two general layouts: object list or details. The recommended display setting is at least 1280 by 1024 pixels.

Not every window contains every element in the following diagrams.

Object list window layout



Object details window layout



Window layout customization


Active IQ Unified Manager enables you to customize the layout of information on the storage and network object pages. By customizing the windows, you can control which data is viewable and how the data is displayed.

• Sorting

You can click the column header to change the sort order of the column entries. When you click the column

header, the sort arrows (▲ and ▼) appear for that column.

- **Filtering**

You can click the filter icon () to apply filters to customize the display of information on the storage and network object pages so that only those entries that match the conditions that are provided are displayed. You apply filters from the Filters pane.

The Filters pane enables you to filter most of the columns based on the options that are selected. For example, on the Health: All Volumes view, you can use the Filters pane to display all of the volumes that are offline by selecting the appropriate filter option under State.

Capacity-related columns in any list always display capacity data in appropriate units rounded off to two decimal points. This also applies when filtering capacity columns. For example, if you use the filter in the Total Data Capacity column in the Health: All Aggregates view to filter data greater than 20.45 GB, the actual capacity of 20.454 GB is displayed as 20.45 GB. Similarly, if you filter data less than 20.45 GB, the actual capacity of 20.449 GB is displayed as 20.45 GB.

If you use the filter in the Available Data % column in the Health: All Aggregates view to filter data greater than 20.45%, the actual capacity of 20.454% is displayed as 20.45%. Similarly, if you filter data less than 20.45%, the actual capacity of 20.449% is displayed as 20.45%.

- **Hiding or showing the columns**

You can click the column display icon (**Show/Hide**) to select which columns you want to display. Once you have selected the appropriate columns you can re-order them by dragging them using your mouse.

- **Searching**

You can use the search box to search for certain object attributes to help refine the list of items in the inventory page. For example, you can enter "cloud" to refine the list of volumes in the volumes inventory page to see all volumes that have the word "cloud" in them.

- **Exporting data**



You can click the **Reports** button (or **Export** button to export data to a comma-separated values (.csv) file, (.pdf) document, or Microsoft Excel (.xlsx) file and use the exported data to build reports.

Using the Unified Manager Help

The Help includes information about all features included in Active IQ Unified Manager. You can use the table of contents, the index, or the search tool to find information about the features and how to use them.

Help is available from each tab and from the menu bar of the Unified Manager user interface.

The search tool in the Help does not work for partial words.

- To learn about specific fields or parameters, click .
- To view all the Help contents, click  > **Help/Documentation** in the menu bar.

You can find more detailed information by expanding any portion of the Table of Contents in the navigation pane.

- To search the Help contents, click the **Search** tab in the navigation pane, type the word or series of words you want to find, and click **Go!**
- To print Help topics, click the printer icon.

Bookmarking your favorite Help topics

In the Help Favorites tab, you can bookmark Help topics that you use frequently. Help bookmarks provide fast access to your favorite topics.

Steps

1. Navigate to the Help topic that you want to add as a favorite.
2. Click **Favorites**, and then click **Add**.

Searching for storage objects

To quickly access a specific object, you can use the **Search all Storage Objects** field at the top of the menu bar. This method of global search across all objects enables you to quickly locate specific objects by type. Search results are sorted by storage object type and you can filter them further by object using the drop-down menu.

What you'll need

- You must have one of the following roles to perform this task: Operator, Application Administrator, or Storage Administrator.
- A valid search must contain at least three characters.

When using the drop-down menu value "All", the global search displays the total number of results found in all object categories; with a maximum of 25 search results for each object category. You can select a specific object type from the drop-down menu to refine the search within a specific object type. In this case the returned list is not restricted to the top 25 objects.

The object types you can search for include:

- Clusters
- Nodes
- Storage VMs
- Aggregates
- Volumes
- Qtrees
- SMB Shares
- NFS Shares
- User or Group Quotas
- LUNs
- NVMe Namespaces
- Initiator Groups

- Initiators
- Consistency Group

Entering a workload name returns the list of workloads under the appropriate Volumes or LUNs category.

You can click any object in the search results to navigate to the Health details page for that object. If there is no direct health page for an object, then the Health page of the parent object is displayed. For example, when searching for a specific LUN, the SVM details page on which the LUN resides is displayed.

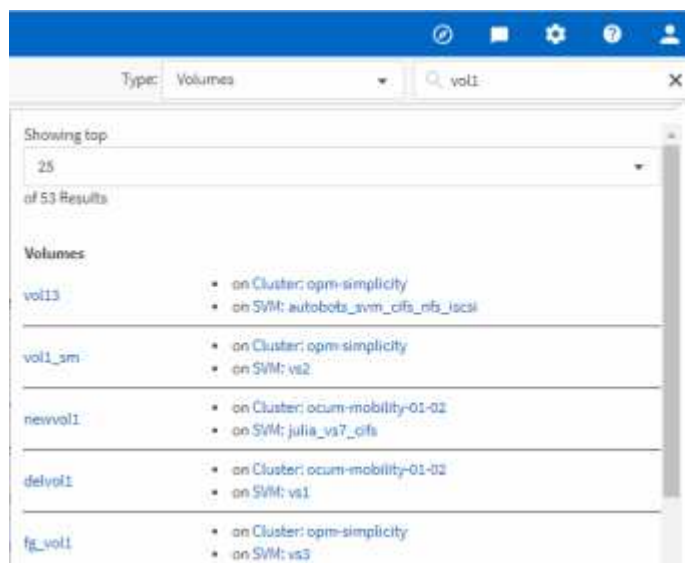


Ports and LIFs are not searchable in the global search bar.

Steps

1. Select an object type from the menu to refine the search results for only a single object type.
2. Type a minimum of three characters of the object name in the **Search all Storage Objects** field.

In this example, the drop-down box has the Volumes object type selected. Typing "vol1" into the **Search all Storage Objects** field displays a list of all volumes whose names contain these characters.



Exporting storage data as reports

You can export storage data in a variety of output formats and then use the exported data to build reports. For example, if there are 10 critical events that have not been resolved, you can export the data from the Event Management inventory page to create a report, and then send the report to admins who can resolve the issues.

You can export data to a .csv file, .xlsx file, or .pdf document from the **Storage** and **Network** inventory pages and use the exported data to build reports. There are other locations in the product where only .csv or .pdf files can be generated.

Steps

1. Perform one of the following actions:

If you want to export...	Do this...
Storage object inventory details	Click Storage or Network from the left-navigation menu, and then select a storage object. Choose one of the system-provided views, or any custom view that you have created.
QoS Policy Group details	Click Storage > QoS Policy Groups from the left-navigation menu.
Storage capacity and protection history details	Click Storage > Aggregates or Storage > Volumes , then select a single aggregate or volume.
Event details	Click Event Management from the left-navigation menu.
Storage object top 10 performance details	Click Storage > Clusters > Performance:All Clusters , then select a cluster and choose the Top Performers tab. Then select a storage object and performance counter.

2. Click the **Reports** button (or **Export** button in some UI pages).
3. Click **Download CSV**, **Download PDF**, or **Download Excel** to confirm the export request.

From the Top Performers tab you can choose to download a report of the statistics for the single cluster you are viewing or for all clusters in the data center.

The file is downloaded.

4. Open the file in the appropriate application.

Related information

[Health/Clusters inventory page](#)

[Scheduling a report](#)

Filtering inventory page content

You can filter inventory page data in Unified Manager to quickly locate data based on specific criteria. You can use filtering to narrow the contents of the Unified Manager pages to show only the results in which you are interested. This provides a very efficient method of displaying only the data in which you are interested.

Use **Filtering** to customize the grid view based on your preferences. Available filter options are based on the object type being viewed in the grid. If filters are currently applied, the number of applied filters displays at the right of the Filter button.

Three types of filter parameters are supported.

Parameter	Validation
String (text)	The operators are contains , starts with , ends with , and does not contain .
Number	The operators are greater than , less than , in the last , and between .
Enum (text)	The operators are is and is not .

The Column, Operator, and Value fields are required for each filter; the available filters reflect the filterable columns on the current page. The maximum number of filters you can apply is four. Filtered results are based on combined filter parameters. Filtered results apply to all pages in your filtered search, not just the page currently displayed.

You can add filters using the Filtering panel.

1. At the top of the page, click the **Filter** button. The Filtering panel displays.
2. Click the left drop-down list and select an object; for example, *Cluster*, or a performance counter.
3. Click the center drop-down list, and select the operator you want to use.
4. In the last list, select or enter a value to complete the filter for that object.
5. To add another filter, click **+Add Filter**. An additional filter field displays. Complete this filter using the process described in the preceding steps. Note that upon adding your fourth filter, the **+Add Filter** button no longer displays.
6. Click **Apply Filter**. The filter options are applied to the grid and the number of filters is displayed to the right of the Filter button.
7. Use the Filtering panel to remove individual filters by clicking the trash icon at the right of the filter to be removed.
8. To remove all filters, click **Reset** at the bottom of the filtering panel.

Filtering example


The illustration shows the Filtering panel with three filters. The **+Add Filter** button displays when you have fewer than the maximum of four filters.

The screenshot shows a Filtering panel with three filter rows. Each row has three main input fields: a column selector (dropdown), an operator selector (dropdown), and a value field. To the right of each value field is a unit or label dropdown and a trash icon for removal. At the bottom left is a '+ Add Filter' button. At the bottom right are 'Cancel' and 'Apply Filter' buttons.


Mbps	greater than	5	Mbps	[trash]
Node	name starts with	test		[trash]
Type	is	FCP Port		[trash]

+ Add Filter

Cancel Apply Filter

After clicking **Apply Filter**, the Filtering panel closes, applies your filters, and shows the number of filters applied ( 3).


Viewing active events from the notification bell

The notification bell () in the menu bar provides a fast way to view the most important active events that Unified Manager is tracking.

The list of active events provides a way to see the total number of critical, error, warning, and upgrade events on all clusters. This list includes events from the previous 7 days, and it does not include Information events. You can click a link to display the list of events that you are most interested in.

Note that when a cluster is not reachable, Unified Manager displays this information in this page. You can view detailed information about a cluster that is unreachable by clicking the **Details** button. This action opens the Event details page. Scale monitoring issues, such as low space or RAM on the management station, are also displayed on this page.

Steps

1. From the menu bar, click .
2. To view details for any of the active events, click the event text link, such as "2 Capacity" or "4 Performance".

Monitoring and managing clusters from the dashboard

The dashboard provides cumulative at-a-glance information about the current health of your monitored ONTAP systems. The dashboard provides “panels” that enable you to assess the overall capacity, performance, and security health of the clusters you are monitoring.

Additionally, there are certain ONTAP issues that you can fix directly from the Unified Manager user interface instead of having to use ONTAP System Manager or the ONTAP CLI.

At the top of the dashboard you can select whether the panels show information for all monitored clusters or for an individual cluster. You can start by viewing the status of all clusters and then drill down to individual clusters when you want to view detailed information.



Some of the panels listed below may not appear on the page based on your configuration.

Panels	Description
Management Actions	When Unified Manager can diagnose and determine a single resolution for an issue, those resolutions are displayed in this panel with a Fix It button.
Capacity	Displays the total and used capacity for the local tier and cloud tier, and the number of days until local capacity reaches the upper limit.
Performance Capacity	Displays the performance capacity value for each cluster and the number of days until performance capacity reaches the upper limit.

Panels	Description
Workload IOPS	Displays the total number workloads that are currently running in a certain range of IOPS.
Workload Performance	Displays the total number of conforming and non-conforming workloads that are assigned to each defined Performance Service Level.
Security	Displays the number of clusters that are compliant or not compliant, the number of SVMs that are compliant or not compliant, and the number of volumes that are encrypted or not encrypted.
Protection	Displays the number of Storage VMs that are protected by SVM-DR relationship, volumes protected by SnapMirror relationship, volumes protected by Snapshot, and clusters protected by MetroCluster.
Usage Overview	Displays clusters sorted by highest IOPS, highest throughput (MBps), or highest used physical capacity.

Dashboard page

The Dashboard page has "panels" that display the high level capacity, performance, and security health of the clusters you are monitoring. This page also provides a Management Actions panel that lists fixes that Unified Manager can make to resolve certain events.

Most of the panels also display the number of active events in that category, and the number of new events added over the previous 24 hours. This information helps you decide which clusters you may need to analyze further to resolve events. Clicking on the events displays the top events and provides a link to the Event Management inventory page filtered to show the active events in that category.

At the top of the dashboard you can select whether the panels show information for all monitored clusters ("All Clusters") or for an individual cluster. You can start by viewing the status of all clusters and then drill down to individual clusters when you want to view detailed information.



Some of the panels listed below appear on the dashboard based on your configuration.

Management Actions panel

There are certain issues that Unified Manager can diagnose thoroughly and provide a single resolution. When available, those resolutions are displayed in this panel with a **Fix It** or **Fix All** button. You can fix these issues immediately from Unified Manager instead of having to use ONTAP System Manager or the ONTAP CLI. For viewing all the issues, click on See [Fixing ONTAP issues directly from Unified Manager](#) for more information.

Capacity panel

When viewing all clusters, this panel displays the physical used capacity (after applying storage efficiency savings) and physical available capacity (not including potential storage efficiency savings) for each cluster, the

number of days until the disks are projected to be full, and the data reduction ratio (without Snapshot copies) based on configured ONTAP storage efficiency settings. It also lists the used capacity for any configured cloud tiers. Clicking the bar chart takes you to the Aggregates inventory page for that cluster. Clicking the "Days To Full" text displays a message that identifies the aggregate with the least number of capacity days remaining; click the aggregate name to see more details.

When viewing a single cluster, this panel displays the physical used capacity and physical available capacity for the data aggregates sorted by each individual disk type on the local tier, and for the cloud tier. Clicking the bar chart for a disk type takes you to the Volumes inventory page for the volumes using that disk type.

Performance Capacity panel

When viewing all clusters, this panel displays the performance capacity value for each cluster (averaged over the previous 1 hour) and the number of days until performance capacity reaches the upper limit (based on daily growth rate). Clicking the bar chart takes you to the Nodes inventory page for that cluster. Note that the Nodes inventory page displays the performance capacity averaged over the previous 72 hours. Clicking the "Days To Full" text displays a message that identifies the node with the least number of performance capacity days remaining; click the node name to see more details.

When viewing a single cluster, this panel displays the cluster performance capacity used percentage, total IOPS, and total throughput (MB/s) values, and the number of days until each of these three metrics are anticipated to reach their upper limit.

Workload IOPS panel

When viewing a single cluster, this panel displays the total number workloads that are currently running in a certain range of IOPS, and indicates the number for each disk type when you hover your cursor over the chart.

Workload Performance panel

This panel displays the total number of conforming and non-conforming workloads that are assigned to each Performance Service Level (PSL) policy. It also displays the number of workloads that are not assigned a PSL. Clicking a bar chart takes you to the conforming workloads assigned to that policy in the Workloads page. Clicking the number that follows the bar chart takes you to the conforming and non-conforming workloads assigned to that policy.

Security panel

The Security panel presents the high-level security status for all clusters or a single cluster, depending on your current view. This panel displays:

- a list of the security events received in the past 24 hours. Click an event to view the details on the Event details page
- the cluster security status (count of compliant and non-compliant clusters)
- the storage VM security status (count of compliant and non compliant storage VMs)
- the volume encryption status (count of the volumes that are encrypted or not encrypted)
- the volume anti-ransomware status (count of volumes with anti-ransomware that are enabled or disabled)

You can click the bar charts of the compliant and non-compliant clusters, storage VMS, encrypted and unencrypted volumes, and volume anti-ransomware status to go to the respective pages and view the security details of the filtered clusters, storage VMs, and volumes.

Compliance is based on the [NetApp Security Hardening Guide for ONTAP 9](#). Click the right-arrow at the top of

the panel to view security details for all clusters on the Security page. For information, see [Viewing detailed security status for clusters and Storage VMs](#).

Data Protection panel

This panel displays the data protection summary for a single or all the clusters in a data center. It displays the total number of data protection events, MetroCluster events, and number of active events raised in the last 24 hours in ONTAP. Clicking the link from each of these events takes you to the Event details page. You can click the **View All** link to view all active protection events in the Event Management inventory page. The panel displays:

- The number of volumes in a cluster or all the clusters in a data center protected by Snapshot copies.
- The number of volumes in a cluster or all the clusters in a data center protected by SnapMirror relationships. For SnapMirror relationships, the volume count at the source cluster is considered.
- The number of clusters or all the clusters in a data center protected by MetroCluster configuration over IP or FC.
- The number of volume relationships with SnapMirror recovery point objective (RPO) lag based on the lag status.

You can hover your mouse to view the respective counts and legends. You can click the right arrow at the top of the panel to view the details for a single or all the clusters on the Data Protection page. Also, you can click:

- The bar charts for unprotected volumes and volumes protected by Snapshot copies to go to the Volumes page and view the details.
- The bar charts for the clusters protected or not protected by MetroCluster configuration to go to the Clusters page and view the details.
- The bar charts for all relationships to go to the Relationships page, where the details are filtered based on the source cluster.

For more information, see [Viewing volume protection status](#).

Usage Overview panel

When viewing all clusters, you can choose to view clusters sorted by highest IOPS, highest throughput (MB/s), or highest used physical capacity.

When viewing a single cluster, you can choose to view workloads sorted by highest IOPS, highest throughput (MB/s), or highest used logical capacity.

Related information

[Fixing issues using Unified Manager automatic remediations](#)

[Displaying information about performance events](#)

[Managing performance using performance capacity and available IOPS information](#)

[Volume / Health details page](#)

[Performance event analysis and notification](#)

[Description of event severity types](#)

[Sources of performance events](#)

[Managing cluster security objectives](#)

[Monitoring cluster performance from the Performance Cluster Landing page](#)

[Monitoring performance using the Performance Inventory pages](#)

Managing ONTAP issues or features directly from Unified Manager

You can fix certain ONTAP issues or manage certain ONTAP features directly from the Unified Manager user interface, instead of having to use ONTAP System Manager or the ONTAP CLI. The “Management Actions” option provides fixes to a number of ONTAP issues that have triggered Unified Manager events.

You can fix issues directly from the Management Actions page by selecting the **Management Actions** option on the left navigation pane. Management Actions are also available from the Management Actions panel on the Dashboard, Event details page, and Workload Analysis selection on the left-navigation menu.

There are certain issues that Unified Manager can diagnose thoroughly and provide a single resolution. For certain ONTAP features, such as anti-ransomware monitoring, Unified Manager performs internal checks and recommends specific actions. When available, those resolutions are displayed in Management Actions with a **Fix It** button. Click the **Fix It** button to fix the issue. You must have the Application Administrator or Storage Administrator role.

Unified Manager sends ONTAP commands to the cluster to make the requested fix. When the fix is complete the event is obsoleted.

Some management actions enable you to fix the same issue on multiple storage objects using the **Fix All** button. For example, there may be 5 volumes that have the "Volume Space Full" event that could be resolved by clicking the **Fix All** management action for "Enable volume autogrow". One click enables you to fix this issue on 5 volumes.

For information about the ONTAP issues and features that you can manage by using automatic remediation, see [What issues can Unified Manager fix](#).

What options do I have when I see the Fix It or Fix All button

The Management Actions page provides you with the **Fix It** or **Fix All** button to fix issues that Unified Manager has been notified about through an event.

We recommend that you click the buttons to fix an issue, as required. However, if you are not sure that you want to resolve the issue as recommended by Unified Manager, you can perform the following actions:

What do you want to do?	Action
Have Unified Manager fix the issue on all identified objects.	Click the Fix All button.
Do not fix the issue for any of the identified objects at this time and hide this management action until the event is raised again.	Click the down arrow and click Dismiss All .

What do you want to do?	Action
Fix the issue on only some of the identified objects.	Click the name of the management action to expand the list and show all individual Fix It actions. Then follow the steps for fixing or dismissing individual management actions.

What do you want to do?	Action
Have Unified Manager fix the issue.	Click the Fix It button.
Do not fix the issue at this time and hide this management action until the event is raised again.	Click the down arrow and click Dismiss .
Display the details for this event so you can better understand the issue.	<ul style="list-style-type: none"> Click the Fix It button and review the fix that will be applied in the resulting dialog box. Click the down arrow and click View Event Detail to display the Event details page. <p>Then click Fix It from either of these resulting pages if you want to fix the issue.</p>
Display the details for this storage object so you can better understand the issue.	Click the name of the storage object to display details in either the Performance Explorer or Health Details page.

In some cases the fix is reflected in the next 15 minute configuration poll. In other cases it can take up to many hours for the configuration change to be verified and for the event to be obsoleted.

To see the list of completed or in progress management actions, click the filter icon and select **Completed** or **In Progress**.

Fix All operations run in a serial fashion, so when you view the **In Progress** panel some objects will have the Status **In Progress** whereas others will have the Status **Scheduled**; meaning they are still waiting to be implemented.

Viewing the status of management actions you have chosen to fix


You can view the status of all management actions that you have chosen to fix in the Management Actions page. Most actions are shown as **Completed** fairly quickly after Unified Manager sends the ONTAP command to the cluster. However, some actions, such as moving a volume, can take longer.

There are three filters available on the Management Actions page:

- **Completed** shows both management actions that completed successfully and those that have failed. **Failed** actions provide a reason for the failure so that you can address the issue manually.
- **In Progress** shows both those management actions that are being implemented, and those that are scheduled to be implemented.

- **Recommended** shows all the management actions that are currently active for all monitored clusters.

Steps

1. Click **Management Actions** on the left navigation pane. Alternately, click  at the top of the **Management Actions** panel on the **Dashboard** and select the View you want to see.

The Management Actions page is displayed.

2. You can click the caret icon next to the management action in the **Description** field to see details about the issue and the command that is being used to fix the issue.
3. To view any actions that **failed**, sort on the **Status** column in the **Completed** View. You can use the **Filter** tool for this same purpose.
4. If you want to view more information about a Failed management action, or if you decide that you want to fix a Recommended management action, you can click **View Event Detail** from the expanded area after you click the caret icon next to the management action. A **Fix It** button is available from that page.

What issues can Unified Manager fix

By using the automatic remediation feature of Active IQ Unified Manager, you can choose to resolve certain ONTAP issues or manage certain ONTAP features, such as anti-ransomware monitoring, effectively through Unified Manager.

This table describes these ONTAP issues or features that you can manage directly through the **Fix It** or **Fix All** button on the Unified Manager web UI.

Event Name and Description	Management Action	"Fix It" Operation
Volume Space Full The volume is almost out of space and it has breached the capacity full threshold. This threshold is set by default to 90% of the volume size.	Enable volume autogrow	Unified Manager determines that volume autogrow is not configured for this volume, so it enables this feature so the volume will grow in capacity when needed.
Inodes Full This volume has run out of inodes and cannot accept any new files.	Increase number of inodes on volume	Increases the number of inodes on the volume by 2 percent.
Storage Tier Policy Mismatch Detected The volume has lots of inactive data and the current tiering policy is set to "snapshot-only" or "none".	Enable automatic cloud tiering	Since the volume already resides on a FabricPool, it changes the tiering policy to "auto" so that inactive data is moved to the lower cost cloud tier.

Event Name and Description	Management Action	"Fix It" Operation
<p>Storage Tier Mismatch Detected</p> <p>The volume has lots of inactive data, but it does not reside on a cloud-enabled storage tier (FabricPool).</p>	Change volumes' storage tier	Moves the volume to cloud-enabled storage tier and sets the tiering policy to "auto" to move inactive data to the cloud tier.
<p>Audit Log Disabled</p> <p>The audit log is not enabled for the storage VM</p>	Enable audit logging for the storage VM	<p>Enables audit logging on the storage VM.</p> <p>Note that the storage VM must already have either a local or remote audit log location configured.</p>
<p>Login Banner Disabled</p> <p>The login banner for the cluster should be enabled to increase security by making access restrictions clear.</p>	Set login banner for the cluster	Sets the cluster login banner to "Access restricted to authorized users".
<p>Login Banner Disabled</p> <p>The login banner for the storage VM should be enabled to increase security by making access restrictions clear.</p>	Set login banner for the storage VM	Sets the storage VM login banner to "Access restricted to authorized users".
<p>SSH is Using Insecure Ciphers</p> <p>Ciphers with the suffix "-cbc" are considered insecure.</p>	Remove insecure ciphers from the cluster	Removes the insecure ciphers — such as aes192-cbc and aes128-cbc — from the cluster.
<p>SSH is Using Insecure Ciphers</p> <p>Ciphers with the suffix "-cbc" are considered insecure.</p>	Remove insecure ciphers from the storage VM	Removes the insecure ciphers — such as aes192-cbc and aes128-cbc — from the storage VM.
<p>AutoSupport HTTPS transport disabled</p> <p>The transport protocol used to send AutoSupport messages to technical support should be encrypted.</p>	Set HTTPS as the transport protocol for AutoSupport messages	Sets HTTPS as the transport protocol for AutoSupport messages on the cluster.

Event Name and Description	Management Action	"Fix It" Operation
<p>Cluster Load Imbalance Threshold Breached</p> <p>Indicates that the load is imbalanced among the nodes in the cluster. This event is generated when the performance capacity used variance is more than 30% between nodes.</p>	Balance cluster workloads	Unified Manager identifies the best volume to move from one node to the other to reduce the imbalance, and then moves the volume.
<p>Cluster Capacity Imbalance Threshold Breached</p> <p>Indicates that the capacity is imbalanced among the aggregates in the cluster. This event is generated when the used capacity variance is more than 70% between aggregates.</p>	Balance cluster capacity	Unified Manager identifies the best volume to move from one aggregate to another to reduce the imbalance, and then moves the volume.
<p>Performance Capacity Used Threshold Breached</p> <p>Indicates that the load on the node could become over utilized if you don't reduce the utilization by one or more highly active workloads. This event is generated when the node performance capacity used value is more than 100% for more than 12 hours.</p>	Limit high load on node	Unified Manager identifies the volume with the highest IOPS and it applies a QoS policy using the historical expected and peak IOPS levels to reduce the load on the node.
<p>Dynamic Event Warning Threshold Breached</p> <p>Indicates that the node is already operating in an overloaded state due to the abnormally high load of some of the workloads.</p>	Reduce overload in node	Unified Manager identifies the volume with the highest IOPS and it applies a QoS policy using the historical expected and peak IOPS levels to reduce the load on the node.
<p>Takeover is not possible</p> <p>Failover is currently disabled, so access to the node's resources during an outage or reboot would be lost until the node became available again.</p>	Enable node failover	Unified Manager sends the appropriate command to enable failover on all nodes in the cluster.

Event Name and Description	Management Action	"Fix It" Operation
<p>Option Cf.takeover.on_panic is Configured OFF</p> <p>The nodeshell option "cf.takeover.on_panic" is set to off, which could cause an issue on HA-configured systems.</p>	Enable takeover on panic	Unified Manager sends the appropriate command to the cluster to change this setting to on .
<p>Disable nodeshell option snapmirror.enable</p> <p>The old nodeshell option "snapmirror.enable" is set to on, which could cause an issue during boot after upgrading to ONTAP 9.3 or greater.</p>	Set snapmirror.enable option to off	Unified Manager sends the appropriate command to the cluster to change this setting to off .
<p>Telnet enabled</p> <p>Indicates a potential security issue because Telnet is insecure and passes data in an unencrypted manner.</p>	Disable Telnet	Unified Manager sends the appropriate command to the cluster to disable Telnet.
<p>Configure storage VM anti-ransomware learning</p> <p>Periodically checks for clusters with licenses for anti-ransomware monitoring. Validates whether a storage VM supports only NFS or SMB volumes in such a cluster.</p>	Put storage VMs in a <code>learning</code> mode of anti-ransomware monitoring	Unified Manager sets anti-ransomware monitoring to <code>learning</code> state for the storage VMs through the cluster management console. Anti-ransomware monitoring on all the new volumes created on the storage VM are automatically moved to a learning mode. Through this enablement, ONTAP can learn the pattern of activity on the volumes and detect the anomalies due to potential malicious attacks.
<p>Configure volume anti-ransomware learning</p> <p>Periodically checks for clusters with licenses for anti-ransomware monitoring. Validates whether a volume supports only NFS or SMB services in such a cluster.</p>	Put volumes in <code>learning</code> mode of anti-ransomware monitoring	Unified Manager sets anti-ransomware monitoring to <code>learning</code> state for the volumes through the cluster management console. Through this enablement, ONTAP can learn the pattern of activity on the volumes and detect the anomalies due to potential malicious attacks.

Event Name and Description	Management Action	"Fix It" Operation
<p>Enable volume anti-ransomware</p> <p>Periodically checks for clusters with licenses for anti-ransomware monitoring. Detects whether the volumes are in the <code>learning</code> mode of anti-ransomware monitoring for more than 45 days, and determines the prospect of putting them in active mode.</p>	Put volumes in <code>active</code> mode of anti-ransomware monitoring	Unified Manager sets anti-ransomware monitoring to <code>active</code> on the volumes through the cluster management console. Through this enablement, ONTAP can learn the pattern of activity on the volumes and detect the anomalies due to potential malicious attacks, and create alerts for data protection actions.
<p>Disable volume anti-ransomware</p> <p>Periodically checks for clusters with licenses for anti-ransomware monitoring. Detects repetitious notifications during active anti-ransomware monitoring on the volumes (for example, multiple warnings of potential ransomware attacks are returned over 30 days).</p>	Disable anti-ransomware monitoring on volumes	Unified Manager disables anti-ransomware monitoring on the volumes through the cluster management console.

Overriding management actions through scripts

You can create custom scripts and associate them to alerts to take specific actions for specific events, and not opt for the default management actions available for them on the Management Actions page or Unified Manager dashboard.

If you want to take specific actions for an event type and choose not to fix them as a part of the management action capability provided by Unified Manager, you can configure a custom script for the specific action. You can then associate the script with an alert for that event type and take care of such events individually. In this case, management actions are not generated for that specific event type on the Management Actions page or Unified Manager dashboard.

Managing clusters

You can manage the ONTAP clusters by using Unified Manager to monitor, add, edit, and remove clusters.

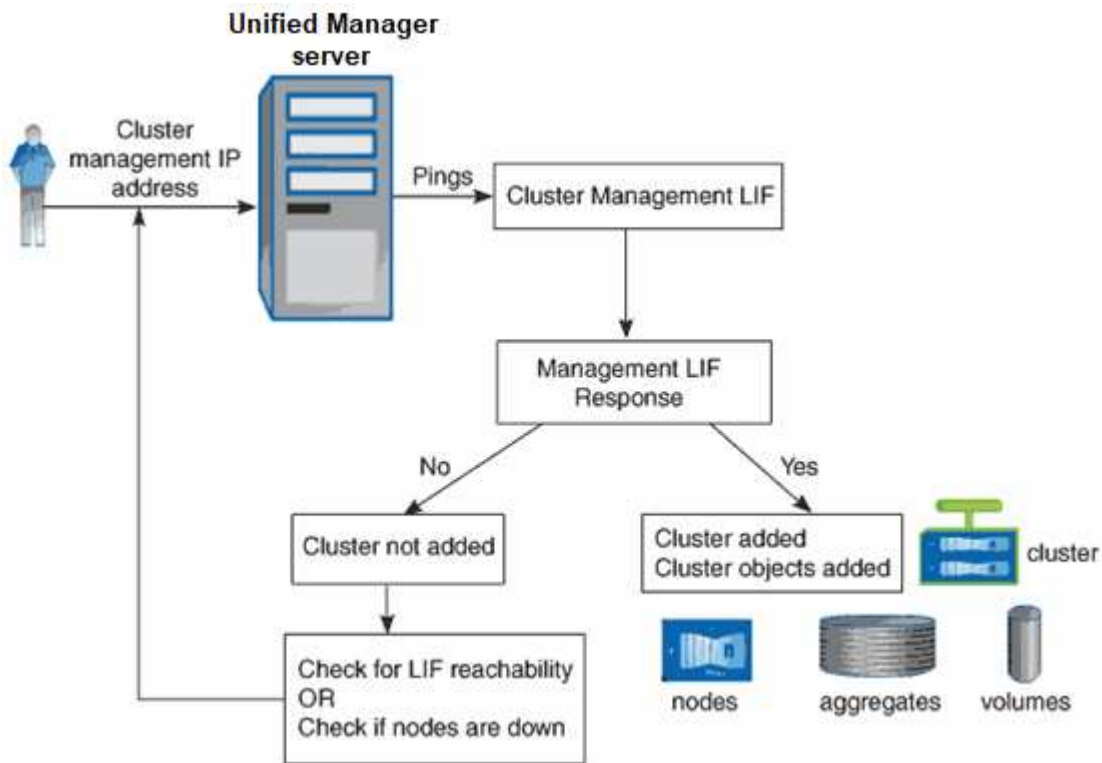
How the cluster discovery process works

After you have added a cluster to Unified Manager, the server discovers the cluster objects and adds them to its database. Understanding how the discovery process works helps you to manage your organization's clusters and their objects.

The monitoring interval for collecting cluster configuration information is 15 minutes. For example, after you have added a cluster, it takes 15 minutes to display the cluster objects in the Unified Manager UI. This time frame is also true when making changes to a cluster. For example, if you add two new volumes to an SVM in a

cluster, you see those new objects in the UI after the next polling interval, which could be up to 15 minutes.

The following image illustrates the discovery process:



After all the objects for a new cluster are discovered, Unified Manager starts to gather historical performance data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.



Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time.

Viewing the list of monitored clusters

You can use the Cluster Setup page to view your inventory of clusters. You can view details about the clusters, such as their name or IP address and communication status.

What you'll need

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

Step

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

All the clusters in your storage environment managed by Unified Manager are displayed. The list of clusters is sorted by the collection state severity level column. You can click a column header to sort the clusters by different columns.

Adding clusters

You can add a cluster to Active IQ Unified Manager so that you can monitor the cluster. This includes the ability to obtain cluster information such as the health, capacity, performance, and configuration of the cluster so that you can find and resolve any issues that might occur.

What you'll need

- You must have the Application Administrator role or the Storage Administrator role.
- You must have the following information:
 - Unified Manager supports on-premises ONTAP clusters, ONTAP Select, Cloud Volumes ONTAP.
 - You must have the host name or cluster management IP address (IPv4 or IPv6) for the cluster.

When using the host name, it must resolve to the cluster management IP address for the cluster management LIF. If you use a node management LIF, the operation fails.

- You must have the user name and password to access the cluster.

This account must have the *admin* role with Application access set to *ontapi*, *console*, and *http*.

- You must know the port number to connect to the cluster using the HTTPS protocol (typically port 443).
- The cluster must be running ONTAP version 9.1 software or greater.
- You must have adequate space on the Unified Manager server. You are prevented from adding a cluster to the server when greater than 90% of space is already consumed.
- You have the required certificates:

SSL (HTTPS) certificate: This certificate is owned by Unified Manager. A default self-signed SSL (HTTPS) certificate is generated with a fresh installation of Unified Manager. NetApp recommends that you upgrade it to CA-signed certificate for better security. If the server certificate expires, you should regenerate it and restart Unified Manager for the services to incorporate the new certificate. For more information about regenerating SSL certificate, see [Generating an HTTPS security certificate](#).

EMS certificate: This certificate is owned by Unified Manager. It is used during authentication for EMS notifications received from ONTAP.

Certificates for Mutual TLS communication: Used during Mutual TLS communication between Unified Manager and ONTAP. The certificate-based authentication is enabled for a cluster, based on the ONTAP version. If the cluster running the ONTAP version is lower than the 9.5, certificate-based authentication is not enabled.

Certificate-based authentication is not enabled automatically for a cluster, if you are updating an older version of Unified Manager. However, you can enable it by modifying and saving the cluster details. If the certificate expires, you should regenerate it to incorporate the new certificate. For more information about viewing and regenerating the certificate, see [Editing clusters](#).



- You can add a cluster from the web UI, and certificate-based authentication is automatically enabled.
- You can add a cluster through Unified Manager CLI, the certificate-based authentication is not enabled by default. If you add a cluster using Unified Manager CLI, it is required to edit the cluster using Unified Manager UI. You can see [Supported Unified Manager CLI commands](#) to add a cluster using Unified Manager CLI.
- If certificate-based authentication is enabled for a cluster, and you take the backup of Unified Manager from a server and restore to another Unified Manager server where hostname or IP address is changed, then monitoring of the cluster can fail. To avoid the failure, edit and save the cluster details. For more information about editing cluster details, see [Editing clusters](#).
- On the cluster level, the Active IQ interface adds two new user group entries for the authentication method "cert".

Cluster certificates: This certificate is owned by ONTAP. You cannot add a cluster to Unified Manager with an expired certificate and if the certificate has already expired, you should regenerate it before adding the cluster. For information about certificate generation, see the knowledge base (KB) article [How to renew an ONTAP self-signed certificate in System Manager user interface](#).

- A single instance of Unified Manager can support a specific number of nodes. If you need to monitor an environment that exceeds the supported node count, you must install an additional instance of Unified Manager to monitor some of the clusters. To view the list of supported node count, see the [Unified Manager Best Practices Guide](#).

Steps

1. In the left navigation pane, click **Storage Management > Cluster Setup**.
2. On the Cluster Setup page, click **Add**.
3. In the Add Cluster dialog box, specify the values as required, and then click **Submit**.
4. In the Authorize Host dialog box, click **View Certificate** to view the certificate information about the cluster.
5. Click **Yes**.

After saving the cluster details, you can see the certificate for Mutual TLS communication for a cluster.

If the certificate-based authentication is not enabled, Unified Manager checks the certificate only when the cluster is added initially. Unified Manager does not check the certificate for each API call to ONTAP.

After all the objects for a new cluster are discovered, Unified Manager starts to gather historical performance data for the previous 15 days. These statistics are collected using the data continuity collection functionality. This feature provides you with over two weeks of performance information for a cluster immediately after it is added. After the data continuity collection cycle is completed, real-time cluster performance data is collected, by default, every five minutes.



Because the collection of 15 days of performance data is CPU intensive, it is suggested that you stagger the addition of new clusters so that data continuity collection polls are not running on too many clusters at the same time. Additionally, if you restart Unified Manager during the data continuity collection period, the collection will be halted and you will see gaps in the performance charts for the missing timeframe.

If you receive an error message that you cannot add the cluster, check to see if the following issues exist:



- If the clocks on the two systems are not synchronized and the Unified Manager HTTPS certificate start date is later than the date on the cluster. You must ensure that the clocks are synchronized using NTP or a similar service.
- If the cluster has reached the maximum number of EMS notification destinations the Unified Manager address cannot be added. By default only 20 EMS notification destinations can be defined on the cluster.

Related information

[Adding users](#)

[Viewing the cluster list and details](#)

[Installing a CA signed and returned HTTPS certificate](#)

Editing clusters

You can modify the settings of an existing cluster, such as the host name or IP address, user name, password, and port, by using the Edit Cluster dialog box.

What you'll need

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



Starting with Unified Manager 9.7, clusters can be added using HTTPS only.

Steps

1. In the left navigation pane, click **Storage Management > Cluster Setup**.
2. On the **Cluster Setup** page, select the cluster you want to edit, and then click **Edit**.
3. In the **Edit Cluster** dialog box, modify the values as required.
If you have modified the details for a cluster added to Unified Manager, you can view the certificate details for Mutual TLS communication, based on the ONTAP version. For more information about ONTAP version, see [Certificates for Mutual TLS communication](#).
You can view the certificate details by clicking **Certificate Details**. If the certificate is expired, click the **Regenerate** button to incorporate the new certificate.
4. Click **Submit**.
5. In the Authorize Host dialog box, click **View Certificate** to view the certificate information about the cluster.
6. Click **Yes**.

Related information

[Adding users](#)

[Viewing the cluster list and details](#)

Removing clusters

You can remove a cluster from Unified Manager by using the Cluster Setup page. For example, you can remove a cluster if cluster discovery fails, or when you want to decommission a storage system.

What you'll need

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

This task removes the selected cluster from Unified Manager. After a cluster is removed, it is no longer monitored. The instance of Unified Manager registered with the removed cluster is also unregistered from the cluster.

Removing a cluster also deletes all its storage objects, historical data, storage services, and all associated events from Unified Manager. These changes are reflected on the inventory pages and the details pages after the next data collection cycle.

Steps

1. In the left navigation pane, click **Storage Management > Cluster Setup**.
2. On the Cluster Setup page, select the cluster that you want to remove and click **Remove**.
3. In the **Remove Data Source** message dialog, click **Remove** to confirm the remove request.

Related information

[Adding users](#)

[Viewing the cluster list and details](#)

Rediscovering clusters

You can manually rediscover a cluster from the Cluster Setup page in order to obtain the latest information about the health, monitoring status, and performance status of the cluster.

You can manually rediscover a cluster when you want to update the cluster—such as by increasing the size of an aggregate when there is insufficient space—and you want Unified Manager to discover the changes that you make.

When Unified Manager is paired with OnCommand Workflow Automation (WFA), the pairing triggers the reacquisition of the data cached by WFA.

Steps

1. In the left navigation pane, click **Storage Management > Cluster Setup**.
2. On the **Cluster Setup** page, click **Rediscover**.

Unified Manager rediscovers the selected cluster and displays the latest health and performance status.

Related information

[Viewing the cluster list and details](#)

Monitoring VMware virtual infrastructure

Active IQ Unified Manager provides visibility into the virtual machines (VMs) in your virtual infrastructure, and enables monitoring and troubleshooting storage and performance issues in your virtual environment. You can use this feature to determine any latency issues in your storage environment or when there is a reported performance event on your vCenter Server.

A typical virtual infrastructure deployment on ONTAP has various components that are spread across compute, network, and storage layers. Any performance lag in a VM application might occur due to a combination of latencies faced by the various components at the respective layers. This feature is useful for storage and vCenter Server admins and IT generalists who need to analyze a performance issue in a virtual environment and understand in which component the issue has occurred.

You can now access the vCenter Server from the vCenter menu of the VMware section. The peek view of each virtual machine listed has the **VCENTER SERVER** link in the TOPOLOGY VIEW that launches the vCenter Server in a new browser. You can also use the **Expand Topology** button to launch the vCenter Server and click **View in vCenter** button to view the datastores in vCenter Server.

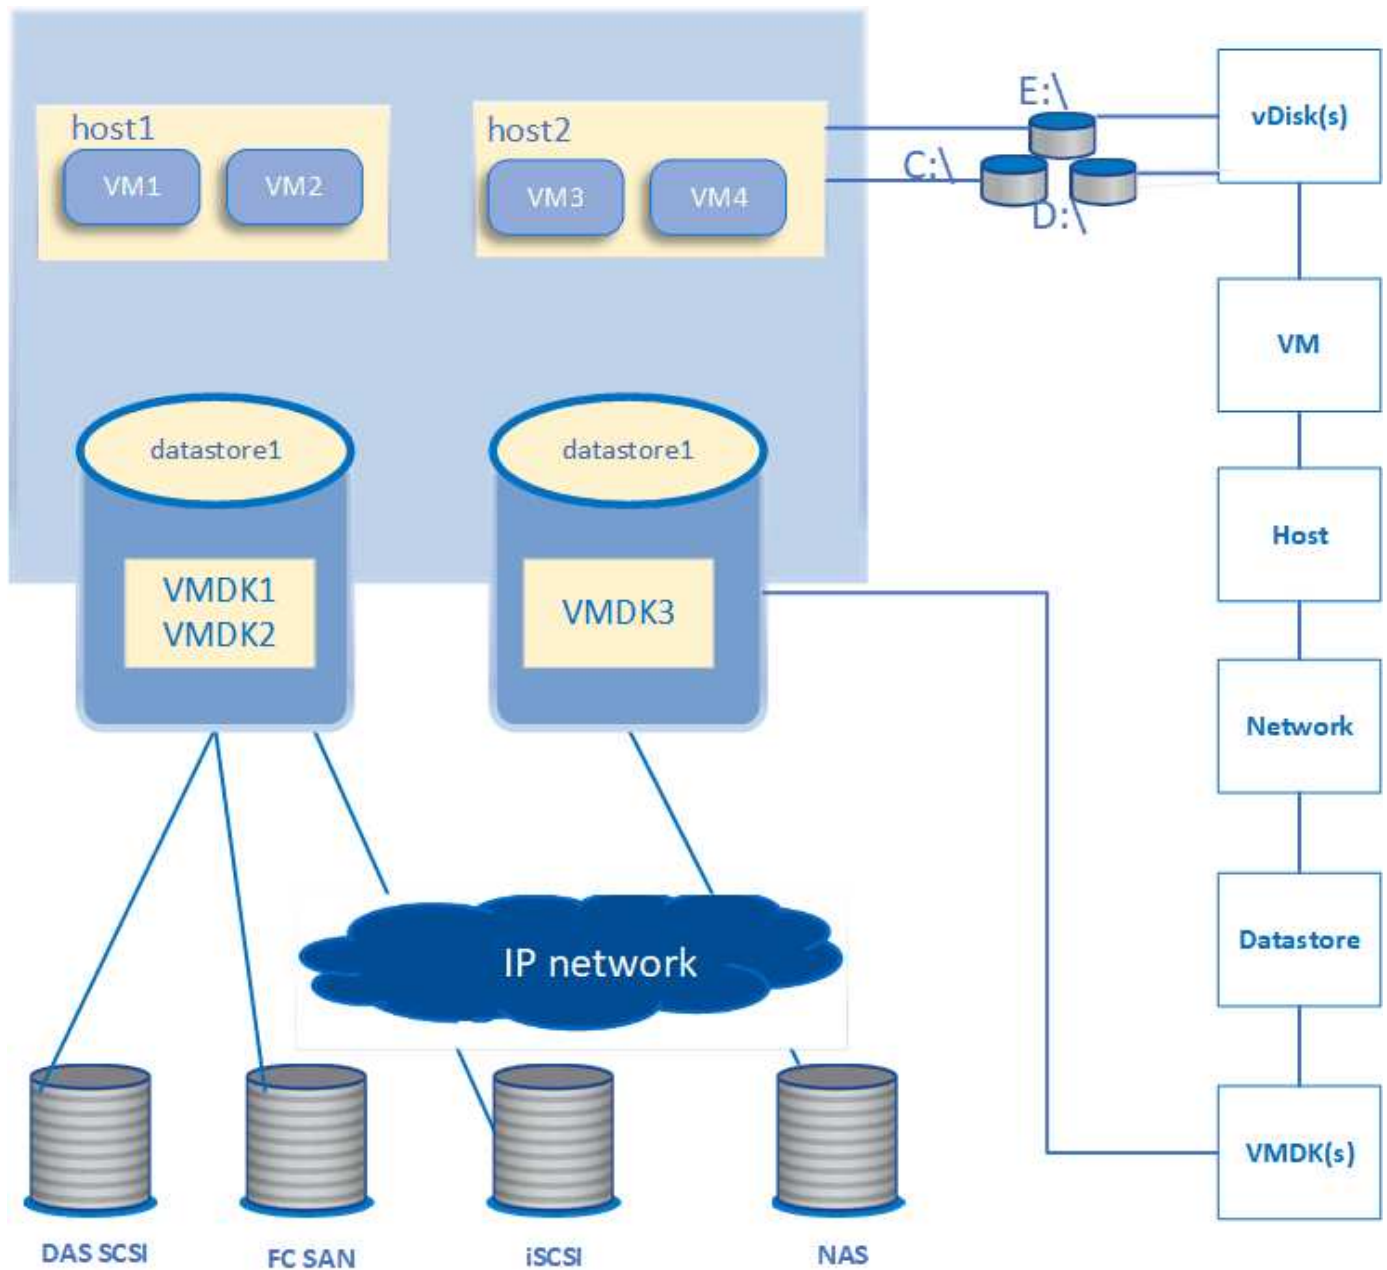
Unified Manager presents the underlying sub-system of a virtual environment in a topological view for determining whether a latency issue has occurred in the compute node, network, or storage. The view also highlights the specific object that causes the performance lag for taking remedial steps and addressing the underlying issue.

A virtual infrastructure deployed on ONTAP storage includes the following objects:

- **vCenter Server:** A centralized control plane for managing the VMware VMs, ESXi hosts, and all related components in a virtual environment. For more information about vCenter Server, see VMware documentation.
- **Host:** A physical or virtual system that runs ESXi, the virtualization software from VMware, and hosts the VM.
- **Datastore:** Datastores are virtual storage objects that are connected to the ESXi hosts. Datastores are manageable storage entities of ONTAP, such as LUNs or volumes, used as a repository for VM files, such as log files, scripts, configuration files, and virtual disks. They are connected to the hosts in the environment via a SAN or IP network connection. Datastores outside of ONTAP that are mapped to vCenter Server are not supported or displayed on Unified Manager.
- **VM:** A VMware virtual machine.
- **Virtual disks:** The virtual disks on datastores belonging to the VMs that have an extension as VMDK. Data from a virtual disk is stored on the corresponding VMDK.
- **VMDK:** A virtual machine disk on the datastore that provides storage space for virtual disks. For each virtual disk, there is a corresponding VMDK.

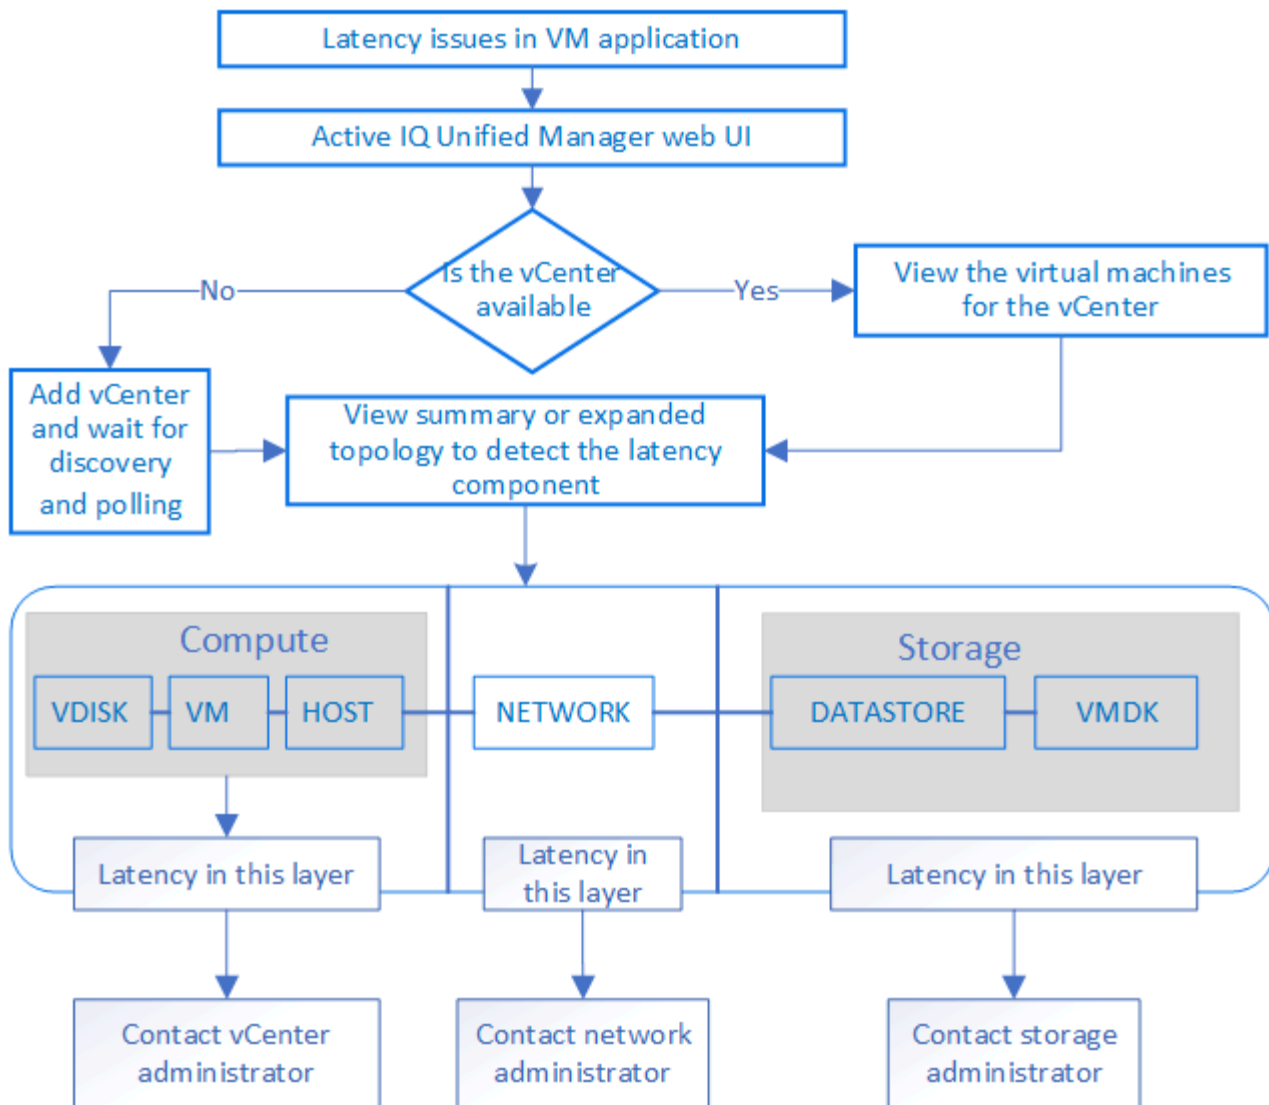
These objects are represented in a VM topology view.

VMware virtualization on ONTAP



User workflow

The following diagram displays a typical use case of using the VM topology view:



What is not supported

- Datastores that are outside of ONTAP and are mapped to the vCenter Server instances are not supported on Unified Manager. Any VMs with virtual disks on those datastores are also not supported.
- A datastore that spans across multiple LUNs is not supported.
- Datastores using Network address translation (NAT) for mapping data LIF (access endpoint) are not supported.
- Exporting volumes or LUNs as datastores on different clusters with same IP addresses in a multiple-LIFs configuration is not supported as Unified Manager is unable to identify which datastore belongs to which cluster.

Example: Suppose cluster A has datastore A. Datastore A is exported via data LIF with same IP address x.x.x.x and VM A is created on this datastore. Similarly, cluster B has datastore B. The datastore B is exported via data LIF with same IP address x.x.x.x and VM B is created on the datastore B. UM will neither be able to map the datastore A for VM A's topology to corresponding ONTAP volume/LUN nor map VM B.

- Only NAS and SAN volumes (iSCSI and FCP for VMFS) are supported as datastores, virtual volumes (vVols) are not supported.
- Only iSCSI virtual disks are supported. Virtual disks of NVMe and SATA types are not supported.

- The views do not allow you to generate reports for analysing the performance of the various components.
- For the storage virtual machine (storage VM) disaster recovery (DR) setup that is supported for only virtual infrastructure on Unified Manager, the configuration has to be manually changed in vCenter Server to point to the active LUNs in switchover and switchback scenarios. Without a manual intervention, their datastores become inaccessible.

Viewing and adding vCenter Server

For viewing and troubleshooting the performance of the virtual machines (VMs), the associated vCenter Servers must be added on your Active IQ Unified Manager instance.

What you'll need

Before adding or viewing vCenter Servers, ensure the following:

- You are aware of the vCenter Server names.
- You know the IP address of vCenter Server and have the required credentials. The credentials must be of a vCenter Server administrator or a root user with read-only access to vCenter Server.
- The vCenter Server that you want to add runs vSphere 6.5 or later.



The support of Unified Manager for VMware ESXi and vCenter Server is available in English and Japanese languages.

- The data collection setting in vCenter Server is set to the statistics level of *Level 3*, ensuring the required level of metrics collection for all the monitored objects. The interval duration should be *5 minutes*, and the save period should be *1 day*.

For more information, see “Data Collection Levels” section of *vSphere Monitoring and Performance Guide* in VMware documentation.

- The latency values in vCenter Server are configured in milliseconds, and not in microseconds, for successful calculations of the latency values.
- While adding the datastore to vCenter Server, you can use both the IP address of the host or the fully qualified domain name (FQDN). In case you are adding FQDN, ensure that the domain name can be resolved by the Unified Manager server. For example, for a Linux installation, ensure that the domain name is added in the `/etc/resolv.conf` file.
- The current time of vCenter Server is in sync with the vCenter Server time zone.
- vCenter Server is reachable for a successful discovery.
- You have the read access to VMware SDK when adding the vCenter Server to Unified Manager. This is required for configuration polling.

For every vCenter Server added and discovered, Unified Manager collects the configuration data, such as the vCenter Server and ESXi server details, ONTAP mapping, datastore details, and number of VMs hosted. It further collects the performance metrics of the components.

Steps

1. Go to **VMWARE > vCenter**, and check whether your vCenter Server is available on the list.



If your vCenter Server is not available, you must add vCenter Server.

- a. Click **Add**.
- b. Add the correct IP address for vCenter Server and ensure that the device is reachable.
- c. Add the user name and password of the administrator or root user with read-only access to vCenter Server.
- d. Add the custom port number if you are using any port other than the default 443.
- e. Click **Save**.

Upon successful discovery, a server certificate is displayed for you to accept.

When you accept the certificate, vCenter Server is added to the list of available vCenter Servers. Adding the device does not result into data collection for the associated VMs, and the collection occurs at the scheduled intervals.

2. If your vCenter Server is available on the **vCenters** page, check its status by hovering your mouse over the **Status** field to display whether your vCenter Server is performing as expected or whether there is a warning or error.



Adding vCenter Server allows you to view the following statuses. However, the performance and latency data of the corresponding VMs might take up to an hour after adding vCenter Server to be accurately reflected.

- Green: "Normal", indicating that vCenter Server has been discovered, and performance metrics have been successfully collected
 - Yellow: "Warning" (for example, when the statistics level for vCenter Server has not been set to 3 or above to obtain statistics for each object)
 - Orange: "Error" (indicates any internal errors, such as exception, failure in configuration data collection, or vCenter Server being unreachable) You can click the column display icon (**Show/Hide**) to view the status message for a vCenter Server status and troubleshoot the issue.
3. In case vCenter Server is unreachable or the credentials have changed, edit the vCenter Server details by selecting **vCenter > Edit**.
 4. Make the necessary changes on the **Edit VMware vCenter Server** page.
 5. Click **Save**.

vCenter Server data collection begins

vCenter Server collects real-time 20-second performance data samples and rolls them up to 5-minute samples. The schedule for performance data collection of Unified Manager is based on the default settings of vCenter Server. Unified Manager processes the 5-minute samples obtained from vCenter Server and calculates an hourly average of the IOPS and latency for the virtual disks, VMs, and hosts. For datastores, Unified Manager calculates an hourly average of the IOPS and latency from samples obtained from ONTAP. These values are available at the top of the hour. The performance metrics are not available immediately after vCenter Server is added, and is available only when the next hour starts. Performance data polling begins on completing a cycle of configuration data collection.

For polling vCenter Server configuration data, Unified Manager follows the same schedule as for collecting cluster configuration data. For information about vCenter Server configuration and performance data collection schedule, see "Cluster configuration and performance data collection activity".

Related information

Removing vCenter Server

You can remove vCenter Servers from your Active IQ Unified Manager instance. For example, you can remove a vCenter Server if vCenter Server discovery fails or when it is no longer required.

Removing a vCenter Server also deletes all the virtual machines (VMs) hosted on that vCenter and its configuration data. After the vCenter Server is removed, it will no longer be monitored, along with its associated objects and historical data. These changes will be reflected on vCenter and virtual machine inventory pages.

What you'll need

Before removing vCenter Servers, ensure the following:

- You have the Application Administrator role or the Storage Administrator role.
- You should be aware of the vCenter Server names and respective IP addresses associated with them.

Steps

1. In the left navigation pane, click **VMWARE>vCenter**.
2. On the vCenters page, select the vCenter Server that you want to remove and click **Remove**.
3. In the **Remove vCenter** message dialog, click **OK** to confirm the remove request.

Monitoring virtual machines

For any latency issue on the virtual machine (VM) applications, you might need to monitor the VMs to analyze and troubleshoot the cause. The VMs are available when their vCenter Server and the ONTAP clusters hosting the VM storage are added to Unified Manager.

You see the details of the VMs on the **VMWARE > > Virtual Machines** page. Information, such as the availability, status, used and allocated capacity, network latency, and the IOPS and latency of the VM, datastore, and host is displayed. For a VM supporting multiple datastores, the grid shows the metrics of the datastore with the worst latency, with an asterisk icon (*) indicating additional datastores. If you click on the icon, the metrics of the additional datastore is displayed. Some of these columns are not available for sorting and filtering.



For viewing a VM and its details, the discovery (polling or metrics collection) of the ONTAP cluster must be complete. If the cluster is removed from Unified Manager, the VM is no longer available, after the next cycle of discovery.

From this page, you can also view the detailed topology of a VM, displaying the components to which the VM is related, for example, the host, virtual disk, and datastore connected to it. The topology view displays the underlying components in their specific layers, in the following order: **Virtual Disk > VM > Host > Network > Datastore > VMDK**.

You can determine the I/O path and component-level latencies from a topological aspect and identify whether storage is the cause of the performance issue. The summary view of the topology displays the I/O path, and highlights the component that has IOPS and latency issues for you to decide on the troubleshooting steps. You

can also have an expanded view of the topology that depicts each component separately along with latency of that component. You can select a component to determine the I/O path highlighted through the layers.

Viewing summary topology

To determine performance issues by viewing the VMs in a summary topology:

1. Go to **VMWARE > Virtual Machines**.
2. Search for your VM by typing its name in the search box. You can also filter your search results based on specific criteria by clicking on the **Filter** button. However, If you cannot find your VM, ensure that the corresponding vCenter Server has been added and discovered.



vCenter Servers allow special characters (such as %, &, *, \$, #, @, !, \, /, :, *, ?, ", <, >, |, ;, ') in the names of vSphere entities, such as VM, cluster, datastore, folder, or file. The VMware vCenter Server and ESX/ESXi Server do not escape special characters used in the display names. However, when the name is processed in Unified Manager, it is displayed differently. For example, a VM named as %\$VC_AIQUM_clone_191124% in vCenter Server is displayed as %25\$VC_AIQUM_clone_191124%25 in Unified Manager. You must keep a note of this issue when you query a VM with a name having a special characters in it.

3. Check the status of the VM. The VM statuses are retrieved from vCenter Server. The following statuses are available. For more information about these statuses, see VMware documentation.
 - Normal
 - Warning
 - Alert
 - Not monitored
 - Unknown
4. Click the down arrow beside the VM to see the summary view of the topology of the components across the compute, network, and storage layers. The node that has latency issues is highlighted. The summary view displays the worst latency of the components. For example, if a VM has more than one virtual disk, this view displays the virtual disk that has the worst latency among all the virtual disks.
5. To analyze the latency and throughput of the datastore over a period of time, click the **Workload Analyzer** button on top of the datastore object icon. You go to the Workload Analysis page, where you can select a time range and view the performance charts of the datastore. For more information about workload analyzer, see *Troubleshooting workloads using the workload analyzer*.

Viewing expanded topology

You can drill down to each component separately by viewing the expanded topology of the VM.

Steps

1. From the summary topology view, click **Expand Topology**. You can see the detailed topology of each component separately with the latency numbers for each object. If there are multiple nodes in a category, for example multiple nodes in the datastore or VMDK, the node with worst latency is highlighted in red.
2. To check the IO path of a specific object, click on that object to see the IO path and the corresponding mapping. For example, to see the mapping of a virtual disk, click the virtual disk to view its highlighted mapping to the respective VMDK. In case of a performance lag of these components, you can collect more data from ONTAP and troubleshoot the issue.



Metrics are not reported for VMDKs. In the topology, only the VMDK names are displayed, and not metrics.

Related information

[Troubleshooting workloads using the workload Analyzer](#)

Viewing virtual infrastructure in a disaster recovery setup

You can view the configuration and performance metrics of the datastores hosted in a MetroCluster configuration or storage virtual machine (storage VM) disaster recovery (SVM DR) setup.

On Unified Manager, you can view the NAS volumes or LUNs in a MetroCluster configuration that are attached as datastores in vCenter Server. The datastores hosted in a MetroCluster configuration are represented in the same topological view as a datastore in a standard environment.

You can also view the NAS volumes or LUNs in a storage VM disaster recovery configuration that are mapped to the datastores in vCenter Server.

Viewing datastores in MetroCluster configuration

Note the following prerequisites before viewing datastores in a MetroCluster configuration:

- In an event of switchover and switchback, the discovery of the primary and secondary clusters of the HA pair, and vCenter Servers should be complete.
- The primary and secondary clusters of the HA pair, and vCenter Servers must be managed by Unified Manager.
- The required setup must be completed on ONTAP and vCenter Server. For information, see ONTAP and vCenter documentation.

[ONTAP 9 Documentation Center](#)

Follow these steps for viewing datastores:

1. On the **VMWARE > Virtual Machines** page, click the VM that hosts the datastore. Click the **Workload Analyzer** or the datastore object link. In the standard scenario when the primary site hosting the volume or LUN is functioning as expected, you can see the vServer cluster details of the primary site.
2. In case of a disaster, and a consecutive switchover to the secondary site, the datastore link points to the performance metrics of the volume or LUN in the secondary cluster. This is reflected after the next cycle of clusters and vServer discovery (acquisition) is complete.
3. After a successful switchback, the datastore link again reflects the performance metrics of the volume or LUN in the primary cluster. This is reflected after the next cycle of clusters and vServer discovery is complete.

Viewing datastores in storage VM disaster recovery configuration

Note the following prerequisites before viewing datastores in a storage VM disaster recovery configuration:

- In an event of switchover and switchback, the discovery of the primary and secondary clusters of the HA pair, and vCenter Servers should be complete.

- Both the source and destination cluster and storage VM peers should be managed by Unified Manager.
- The required setup must be completed on ONTAP and vCenter Server.
 - For NAS (NFS and VMFS) datastores, in case of a disaster, the steps include bringing up the secondary storage VM, verifying the data LIFs and routes, establishing lost connections on vCenter Server, and starting the VMs.

For a switchback to the primary site, the data between the volumes should be synced before the primary site starts serving the data.

- For SAN (iSCSI and FC for VMFS) datastores, vCenter Server formats the mounted LUN in a VMFS format. In case of a disaster, the steps include bringing up the secondary storage VM, verifying the data LIFs and routes. If the iSCSI target IPs are different from the primary LIFs, they need to be manually added. The new LUNs should be available as devices under the iSCSI adapter of the storage adapter of the host. Thereafter, new VMFS datastores with the new LUNs should be created and the old VMs registered with new names. The VMs must be up and running.

In case of a recovery, the data between the volumes should be synced. New VMFS datastores should again be created using the LUNs and the old VMs registered with new names.

For information about the setup, see ONTAP and vCenter Server documentation.

[ONTAP 9 Documentation Center](#)

Follow these steps for viewing datastores:

1. On the **VMWARE > Virtual Machines** page, click the VM inventory that hosts the datastore. Click the datastore object link. In the standard scenario, you can see the performance data of the volumes and LUNs in the primary storage VM.
2. In case of a disaster, and a consecutive switchover to the secondary storage VM, the datastore link points to the performance metrics of the volume or LUN in the secondary storage VM. This is reflected after the next cycle of clusters and vServer discovery (acquisition) is complete.
3. After a successful switchback, the datastore link again reflects the performance metrics of the volume or LUN in the primary storage VM. This is reflected after the next cycle of clusters and vServer discovery is complete.

Unsupported scenarios

- For a MetroCluster configuration, note the following limitations:
 - Clusters in only the **NORMAL** and **SWITCHOVER** states are taken up. Other states, such as **PARTIAL_SWITCHOVER**, **PARTIAL_SWITCHBACK**, and **NOT_REACHABLE** are not supported.
 - Unless Automatic Switch Over (ASO) is enabled, if the primary cluster goes down, the secondary cluster cannot be discovered, and the topology continues to point to the volume or LUN in the primary cluster.
- For a storage VM disaster recovery configuration, note the following limitation:
 - A configuration with Site Recovery Manager (SRM) or Storage Replication Adapter (SRA) enabled for a SAN storage environment is not supported.

Provisioning and managing workloads

The active management feature of Active IQ Unified Manager provides Performance Service Levels, Storage Efficiency Policies, and storage provider APIs for provisioning, monitoring, and managing storage workloads in a data center.



Unified Manager provides this functionality by default. You can disable it from **Storage Management > Feature Settings** if you do not plan to use this functionality.

When enabled, you can provision workloads on the ONTAP clusters managed by your instance of Unified Manager. You can also assign policies, such as Performance Service Levels and Storage Efficiency Policies on the workloads and manage your storage environment based on those policies.

This feature enables the following functions:

- Automatic discovery of storage workloads on the added clusters enabling easy storage workload evaluation and deployment
- Provisioning NAS workloads supporting NFS and CIFS protocols
- Provisioning SAN workloads supporting iSCSI and FCP protocols
- Support for both NFS and CIFS protocols on the same file share
- Management of Performance Service Levels and Storage Efficiency Policies
- Assigning Performance Service Levels and Storage Efficiency Policies to storage workloads

The **Provisioning**, **Storage > Workloads**, and **Policies** options on the left pane of the UI enable you to modify various configurations.

You can perform the following functions by using these options:

- View storage workloads on the **Storage > Workloads** page
- Create storage workloads from the Provision Workload page
- Create and manage Performance Service Levels from Policies
- Create and manage Storage Efficiency Policies from Policies
- Assign policies to storage workloads from the Workloads page

Related information

[Policy-based storage management](#)

Workloads overview

A workload represents the input/output (I/O) operations of a storage object, such as a volume or LUN. The way the storage is provisioned is based on the expected workload requirements. Workload statistics are tracked by Active IQ Unified Manager only after there is traffic to and from the storage object. For example, the workload IOPS and latency values are available after users start using a database or email application.

The Workloads page displays a summary of the storage workloads of the ONTAP clusters managed by Unified Manager. It provides cumulative at-a-glance information about the storage workloads that conform to the

Performance Service Level, as well as the non-conforming storage workloads. It also enables you to assess the total, available, and used capacity and performance (IOPS) of the clusters across your data center.



It is recommended that you assess the number of storage workloads that are non-conforming, unavailable, or not managed by any Performance Service Level, and take the necessary actions to ensure their conformance, capacity usage, and IOPS.

The Workloads page has the following two sections:

- **Workloads overview:** Provides an overview of the number of storage workloads on the ONTAP clusters managed by Unified Manager.
- **Data center overview:** Provides an overview of the capacity and IOPS of the storage workloads in the data center. The relevant data is displayed at a data center level and for individual .

Workloads overview section

The workloads overview section provides cumulative at-a-glance information of the storage workloads. The status of the storage workloads is displayed based on assigned and unassigned Performance Service Levels.

- **Assigned:** The following statuses are reported for storage workloads on which Performance Service Levels have been assigned:
 - **Conforming:** Performance of storage workloads is based on the Performance Service Levels assigned to them. If the storage workloads are within the threshold latency defined in the associated Performance Service Levels, they are marked as “conforming”. The conforming workloads are marked in blue.
 - **Non-conforming:** During performance monitoring, storage workloads are marked as “non-conforming” if the storage workloads latency exceeds the threshold latency defined in the associated Performance Service Level. The non-conforming workloads are marked in orange.
 - **Unavailable:** Storage workloads are marked as “unavailable” if they are offline, or if the corresponding cluster is unreachable. The unavailable workloads are marked in red.
- **Unassigned:** Storage workloads that do not have a Performance Service Level assigned to them, are reported as “unassigned”. The number is conveyed by the information icon.

The total workload count is the sum total of the assigned and unassigned workloads.

You can click the total number of workloads displayed in this section, and view them on the Workloads page.

The Conformance by Performance Service Levels subsection displays the total number of available storage workloads:

- Conforming to each type of Performance Service Level
- For which there is a mismatch between the assigned and the recommended Performance Service Levels

Data center overview section

The data center overview section graphically represents the available and used capacity, and IOPS for all of the clusters in the data center. By using this data, you should manage the capacity and IOPS of the storage workloads. The section also displays the following information for the storage workloads across all of the clusters:

- The total, available, and used capacity for all of the clusters in your data center

- The total, available, and used IOPS for all of the clusters in your data center
- The available and used capacity based on each Performance Service Level
- The available and used IOPS based on each Performance Service Level
- The total space and IOPS used by the workloads that have no Performance Service Level assigned

How data center capacity and performance is calculated based on Performance Service Levels

The used capacity and IOPS is retrieved in terms of the total used capacity and performance of all of the storage workloads in the clusters.

The available IOPS is calculated based on the expected latency and recommended Performance Service Levels on the nodes. It includes the available IOPS for all of the Performance Service Levels whose expected latency is less than or equal to their own expected latency.

The available capacity is calculated based on the expected latency and recommended Performance Service Levels on aggregates. It includes the available capacity for all of the Performance Service Levels whose expected latency is less than or equal to their own expected latency.

Viewing workloads

When you add clusters to Unified Manager, the storage workloads on each cluster are automatically discovered and displayed on the Workloads page.

Unified Manager begins to analyze the workloads for recommendation (recommended PSLs) only after I/O operations start on the storage workloads.

FlexGroup volumes and its constituents are excluded.

Workloads overview

The Workloads Overview page displays the overview of the workloads in the data center and the space and performance overview of the data center.

- **Workloads Overview** panel: Displays the total number of workloads and the number of workloads with or without PSLs assigned on them. The break-up of the workload count for each PSL is also displayed. Clicking the counts takes you to the **All Workloads** view with the filtered workloads. You can also view the number of workloads that do not conform with the system recommendation and assign the system-recommended PSLs to them by clicking the **Assign System-Recommended PSLs** button.
- **Data Center Overview** panel: Displays the available and used space (TiB) and performance (IOPS) of the data center. A break-up of the available and used space (TiB) and performance (IOPS) of all the workloads under each PSL is also displayed.

All workloads view

The **Storage > Workloads > All Workloads** page lists the storage workloads associated with the ONTAP clusters managed by Unified Manager.

For the newly-discovered storage workloads on which there have been no I/O operations, the status is “Waiting for I/O”. After I/O operations begin on the storage workloads, Unified Manager begins the analysis and the workload status changes to “Learning...”. After the analysis is complete (within 24 hours from the beginning of the I/O operations), the recommended PSLs are displayed for the storage workloads.

The page also enables you to assign Storage Efficiency Policies (SEPs) and Performance Service Levels

(PSLs) to storage workloads. You can perform multiple tasks:

- Add or provision storage workloads
- View and filter the list of workloads
- Assign PSLs to storage workloads
- Evaluate system-recommended PSLs and assign them to workloads
- Assign SEPs to storage workloads

Adding or provisioning storage workloads

You can add or provision the storage workloads to supported LUNs (supporting both iSCSI and FCP protocols), NFS file shares, and SMB shares.

Steps

1. Click **Storage > Workloads > All Workloads > Create**.
2. Create workloads. For information, see [Provisioning and managing workloads](#).

Viewing and filtering workloads

On the All Workloads screen, you can view all the workloads in your data center or search for specific storage workloads based on either their PSLs or names. You can use the filter icon to enter specific conditions for your search. You can search by different filter conditions, such as by the host cluster or storage VM. The **Capacity Total** option allows filtering by the total capacity of the workloads (by MB). However, in this case, the number of workloads returned might vary, because the total capacity is compared at a byte level.

For each workload, information, such as the host cluster and storage VM is displayed, along with the assigned PSL and SEP.

The page also enables you to view the performance details of a workload. You can view detailed information about the IOPS, capacity, and latency of the workload by clicking the **Choose / Order Columns** button and selecting specific columns to view. The Performance View column displays the average and peak IOPS for a workload, and you can click the workload analyser icon to view the detailed IOPS analysis.

Analyzing performance and capacity criteria for a workload

The **Analyze Workload** button on the **IOPS Analysis** pop-up takes you to the Workload Analysis page, where you can select a time range and view the latency, throughput, and capacity trends for the selected workload. For more information about workload analyzer, see [Troubleshooting workloads using the workload analyzer](#).

You can view performance information about a workload to help with troubleshooting by clicking the bar chart icon in the **Performance View** column. To view performance and capacity charts on the Workload Analysis page to analyze the object, click the **Analyze Workload** button.

For more information, see [What data does the workload analyzer display](#).

Assigning policies to workloads

You can assign Storage Efficiency Policies (SEPs) and Performance Service Levels (PSLs) to storage workloads from the All Workloads page by using the different navigation options.

Assigning policies to a single workload

You can assign either a PSL or an SEP or both, to a single workload. Follow these steps:

1. Select the workload.
2. Click the edit icon next to the row, and then click **Edit**.

The **Assigned Performance Service Level** and **Storage Efficiency Policy** fields are enabled.

3. Select the required PSL or SEP, or both.
4. Click the check icon to apply the changes.



You can also select a workload and click **More Actions** to assign the policies.

Assigning policies to multiple storage workloads

You can assign a PSL or an SEP to multiple storage workloads together. Follow these steps:

1. Select the check boxes for the workloads to which you want to assign the policy, or select all the workloads in your data center.
2. Click **More Actions**.
3. For assigning a PSL, select **Assign Performance Service Level**. For assigning an SEP, select **Assign Storage Efficiency Policy**. A pop-up is displayed for you to select the policy.
4. Select the appropriate policy and click **Apply**. The number of workloads on which the policies are assigned are displayed. The workloads on which the policies are not assigned are also listed, with the cause.



Applying policies on workloads in bulk might take some time depending on the number of workloads selected. You can click the **Run in background** button and continue with other tasks while the operation runs in the background. When the bulk assignment is complete, you can view the completion status. If you are applying a PSL on multiple workloads, you cannot trigger another request when the previous job of bulk assignment is running.

Assigning system-recommended PSLs to workloads

You can assign system-recommended PSLs to those storage workloads in a data center that have no PSLs assigned, or the assigned PSLs do not match the system recommendation. To use this functionality, click the **Assign System Recommended PSLs** button. You do not have to select specific workloads.

The recommendation is internally determined by system analytics, and is skipped for those workloads whose IOPS and other parameters do not coincide with the definitions of any available PSL. Storage workloads with *Waiting for I/O* and *Learning* statuses are also excluded.



There are special keywords that Unified Manager looks for in the workload name to override the system analytics and recommend a different PSL for the workload. When the workload has the letters “ora” in the name, the **Extreme Performance** PSL is recommended. And when the workload has the letters “vm” in the name, the **Performance** PSL is recommended.

Also see the knowledge base (KB) article [ActiveIQ Unified Manager 'Assign System Recommended Performance Service Level' is not adaptive to a highly variable workload](#)

Provisioning file share volumes

You can create file share volumes that support CIFS/SMB and NFS protocols, on an existing cluster and Storage Virtual Machine (storage VM) from the Provision Workload page.

What you'll need

- The storage VM must have space for provisioning the file share volume.
- Either or both of the SMB and NFS services should be enabled on your storage VM.
- For selecting and assigning the Performance Service Level (PSL) and Storage Efficiency Policy (SEP) on the workload, the policies must have been created before you start creating the workload.

Steps

1. On the **Provision Workload** page, add the name of the workload that you want to create, and then select the cluster from the available list.
2. Based on the cluster that you have selected, the **STORAGE VM** field filters the available storage VMs for that cluster. Select the required storage VM from the list.

Based on the SMB and NFS services supported on the storage VM, the NAS option is enabled in the Host Information section.

3. In the Storage and Optimization section, assign the storage capacity and PSL, and optionally, an SEP for the workload.

The specifications for the SEP are assigned to the LUN and the definitions for the PSL are applied to the workload when it is created.

4. Select the **Enforce performance limits** check box if you want to enforce the PSL that you have assigned to the workload.

Assigning a PSL to a workload ensures that the aggregate on which the workload is created can support the performance and capacity objectives defined in the respective policy. For example, if a workload is assigned "Extreme Performance" PSL, the aggregate on which the workload is to be provisioned should have the capability of supporting the performance and capacity objectives of the "Extreme Performance" policy, such as SSD storage.



Unless you select this check box, the PSL is not applied to the workload, and the status of the workload on the dashboard appears as unassigned.

5. Select the **NAS** option.

If you cannot see the **NAS** option enabled, verify whether the storage VM that you have selected supports either SMB or NFS, or both.



If your storage VM is enabled for both SMB and NFS services, you can select the **Share by NFS** and **Share by SMB** check boxes and create a file share that supports both NFS and SMB protocols. If you want to create either an SMB or a CIFS share, select only the respective check box.

6. For NFS file share volumes, specify the IP address of the host or network to access the file share volume. You can enter comma-separated values for multiple hosts.

On adding the host IP address, an internal check runs for matching the host details with the storage VM and the export policy for that host is created, or in case there is an existing policy, it is reused. If there are several NFS shares created for the same host, then an available export policy for the same host with matching rules is reused for all the files shares. The function of specifying rules of individual policies or reusing policies by providing specific policy keys is available when you provision the NFS share by using APIs.

7. For an SMB share, specify which users or user groups can access the SMB share and assign the required permissions. For each group of users, a new access control list (ACL) is generated during the file share creation.
8. Click **Save**.

The workload is added to the list of storage workloads.

Provisioning LUNs

You can create LUNs that support CIFS/SMB and NFS protocols, on an existing cluster and Storage Virtual Machine (storage VM) from the Provision Workload page.

What you'll need

- The storage VM must have space for provisioning the LUN.
- Both iSCSI and FCP must be enabled on the storage VM on which you create the LUN.
- For selecting and assigning the Performance Service Level (PSL) and Storage Efficiency Policy (SEP) on the workload, the policies must have been created before you start creating the workload.

Steps

1. On the **Provision Workload** page, add the name of the workload that you want to create, and then select the cluster from the available list.

Based on the cluster that you have selected, the **STORAGE VM** field filters the available storage VMs for that cluster.

2. Select the storage VM from the list that supports the iSCSI and FCP services.

Based on your selection, the SAN option is enabled in the Host Information section.

3. In the **Storage and Optimization** section, assign the storage capacity and PSL, and optionally, the SEP for the workload.

The specifications for the SEP are assigned to the LUN and the definitions for the PSL are applied to the workload when it is created.

4. Select the **Enforce performance limits** check box if you want to enforce the assigned PSL on the workload.

Assigning a PSL to a workload ensures that the aggregate on which the workload is created can support the performance and capacity objectives defined in the respective policy. For example, if a workload is assigned the "Extreme Performance" PSL, the aggregate on which the workload is to be provisioned should have the capability of supporting the performance and capacity objectives of the "Extreme Performance" policy, such as SSD storage.



Unless you select this check box, the PSL is not applied to the workload, and the status of the workload on the dashboard appears as `unassigned`.

5. Select the **SAN** option. If you cannot see the **SAN** option enabled, verify whether the storage VM that you have selected supports iSCSI and FCP.
6. Select the host OS.
7. Specify the host mapping to control access of the initiators to the LUN. You can assign existing initiator groups (igroups), or define and map new igroups.



If you create a new igroup while provisioning the LUN, you need to wait till the next discovery cycle (up to 15 minutes) for using it. It is therefore recommended that you use an existing igroup from the list of available igroups.

If you want to create a new igroup, select the **Create a new initiator group** button, and enter the information for the igroup.

8. Click **Save**.

The LUN is added to the list of storage workloads.

Performance Service Levels

A Performance Service Level (PSL) enables you to define the performance and storage objectives for a workload. You can assign a PSL to a workload when initially creating the workload, or afterwards by editing the workload.

The management and monitoring of storage resources are based on Service Level Objectives (SLOs). SLOs are defined by service level agreements that are based on required performance and capacity. In Unified Manager, SLOs refer to the PSL definitions of the applications that are running on NetApp storage. Storage services are differentiated based on the performance and utilization of the underlying resources. A PSL is a description of the storage service objectives. A PSL enables the storage provider to specify the performance and capacity objectives for the workload. When you assign a PSL on a workload, the corresponding workload on ONTAP is managed by its performance and capacity objectives. Each PSL is governed by peak, expected, and absolute minimum IOPs, and expected latency.

Unified Manager has the following types of PSLs:

- **System-defined:** Unified Manager provides a few canned policies that cannot be changed. These predefined PSLs are:
 - Extreme Performance
 - Performance
 - Value

The Extreme Performance, Performance, and Value PSLs are applicable for most of the common storage workloads in a data center.

Unified Manager also offers three Performance Service Levels for database applications. These are extremely high-performance PSLs that support bursty IOPS and are appropriate for database applications with the highest throughput demand.

- Extreme for Database Logs
- Extreme for Database Shared Data
- Extreme for Database Data
- **User-defined:** If the predefined Performance Service Levels do not meet your requirements, then you can create new PSLs to meet your needs. For information, see [Creating and editing Performance Service Levels](#).
- **Beyond Extreme:** The Beyond Extreme PSLs are the system-recommended PSLs that are suggested for workloads that demand IOPs higher than Extreme. The workloads are internally analyzed based on their IOPS, capacity, and latency, and a Beyond Extreme PSL is recommended for each of these workloads on the **Storage > Workloads > All Workloads** screen. You can apply the PSLs to the workloads to ensure an optimum performance.

The IOPs parameters for the workloads are dynamically generated, depending on the workload behavior, and appended to the name of the Beyond Extreme PSL in the format `Beyond Extreme <number-(peak IOPS/TB)> <number(expected IOPS/TB)>`. For example, if the system determines a workload to have the peak and expected IOPs as 106345 and 37929 respectively, the Beyond Extreme PSL that is generated for the workload is named as `Beyond Extreme 106345 37929`. Though these PSLs are recommended by the system, when you assign them to workloads, these PSLs are labeled as User-defined in type.

Managing workloads by assigning PSLs

You can access PSLs from the **Policies > Performance Service Levels** page and by using the storage provider APIs. Managing storage workloads by assigning PSLs to them is convenient as you do not have to individually manage the storage workloads. Any modifications can also be managed by reassigning another PSL rather than managing them individually. Unified Manager helps you to assign PSLs on your workloads based on internal assessment and recommendations.

For information on assigning system-recommended PSLs to workloads, see [Assigning system-recommended PSLs to workloads](#)

The Performance Service Levels page lists the available PSL policies and enables you to add, edit, and delete them.



You cannot modify a PSL that is system-defined or that is currently assigned to a workload. You cannot delete a PSL that is assigned to a workload, or if it is the only available PSL.

This page displays the following information:

Field	Description
Name	Name of the PSL.
Type	Whether the policy is system-defined or user-defined.
Expected IOPS/TB	Minimum number of IOPS that an application is expected to perform on a LUN or file share. Expected IOPS specifies the minimum expected IOPS allocated, based on the storage object allocated size.

Field	Description
Peak IOPS/TB	<p>Maximum number of IOPS that an application can perform on a LUN or file share. Peak IOPS specifies the maximum possible IOPS allocated, based on the storage object allocated size or the storage object used size.</p> <p>Peak IOPS are based on an allocation policy. The allocation policy is either allocated-space or used-space. When the allocation policy is set to allocated-space, the peak IOPS is calculated based on the size of the storage object. When the allocation policy is set to used-space, the peak IOPS is calculated based on the amount of data stored in the storage object, taking into account storage efficiencies. By default, the allocation policy is set to used-space.</p>
Absolute minimum IOPS	<p>The absolute minimum IOPS is used as an override, when the expected IOPS is less than this value. The default values of the system-defined PSLs are the following:</p> <ul style="list-style-type: none"> • Extreme Performance: If expected IOPS \geq 6144/TB, then absolute minimum IOPS = 1000 • Performance: If expected IOPS \geq 2048/TB and $<$ 6144/TB, then absolute minimum IOPS = 500 • Value: If expected IOPS \geq 128/TB and $<$ 2048/TB, then absolute minimum IOPS = 75 <p>The default values of the system-defined database PSLs are the following:</p> <ul style="list-style-type: none"> • Extreme for Database Logs: If expected IOPS \geq 22528, then absolute minimum IOPS = 4000 • Extreme for Database Shared Data: If expected IOPS \geq 16384, then absolute minimum IOPS = 2000 • Extreme for Database Data: If expected IOPS \geq 12288, then absolute minimum IOPS = 2000 <p>The higher value of the absolute minimum IOPS for custom PSLs can be a maximum of 75000. The lower value is calculated as the following:</p> <p>$1000/\text{expected latency}$</p>
Expected latency	Expected latency for storage IOPS in milliseconds per operation (ms/op).
Capacity	Total available and used capacity in the clusters.

Field	Description
Workloads	Number of storage workloads that have been assigned the PSL.

For information about how the peak IOPS and expected IOPs help in achieving consistent differentiated performance on ONTAP clusters, see the following KB article: [What is Performance Budgeting?](#)

Events generated for workloads breaching the threshold defined by PSLs

Note that if workloads exceed the expected latency value for 30% of the time during the previous hour, Unified Manager generates one of the following events to notify you of a potential performance issue:

- Workload Volume Latency Threshold Breached as defined by Performance Service Level Policy
- Workload LUN Latency Threshold Breached as defined by Performance Service Level Policy.

You may want to analyze the workload to see what may be causing the higher latency values.

For more information, see the following links:

- [Volume events](#)
- [What happens when a performance threshold policy is breached](#)
- [How Unified Manager uses workload latency to identify performance issues](#)
- [What performance events are](#)

System-defined PSLs

The following table provides information about the system-defined PSLs:

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Extreme Performance	Provides extremely high throughput at a very low latency Ideal for latency-sensitive applications	1	12288	6144	1000
Performance	Provides high throughput at a low latency Ideal for database and virtualized applications	2	4096	2048	500

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Value	<p>Provides high storage capacity and moderate latency</p> <p>Ideal for high-capacity applications such as email, web content, file shares, and backup targets</p>	17	512	128	75
Extreme for Database Logs	<p>Provides maximum throughput at the lowest latency.</p> <p>Ideal for database applications supporting database logs. This PSL provides the highest throughput because database logs are extremely bursty and logging is constantly in demand.</p>	1	45056	22528	4000
Extreme for Database Shared Data	<p>Provides very high throughput at the lowest latency.</p> <p>Ideal for database applications data that is stored in a common data store, but is shared across databases.</p>	1	32768	16384	2000

Performance Service Level	Description and use case	Expected latency (ms/op)	Peak IOPS	Expected IOPS	Absolute minimum IOPS
Extreme for Database Data	Provides high throughput at the lowest latency. Ideal for database applications data, such as database table information and metadata.	1	24576	12288	2000

Creating and editing Performance Service Levels

When the system-defined Performance Service Levels do not match your workload requirements, you can create your own Performance Service Levels that are optimized for your workloads.

What you'll need

- You must have the Application Administrator role.
- The Performance Service Level name must be unique, and you cannot use the following reserved keywords:

Prime, Extreme, Performance, Value, Unassigned, Learning, Idle, Default, and None.

You create and edit custom Performance Service Levels from the Performance Service Levels page by defining the service level objectives you require for the applications that will access storage.



You cannot modify a Performance Service Level if it is currently assigned to a workload.

Steps

1. In the left navigation pane under **Settings**, select **Policies > Performance Service Levels**.
2. In the **Performance Service Levels** page, click the appropriate button depending on whether you want to create a new Performance Service Level or if you want to edit an existing Performance Service Level.

To...	Follow these steps...
Create a new Performance Service Level	Click Add .
Edit an existing Performance Service Level	Select an existing Performance Service Level, and then click Edit .

The page to add or edit a Performance Service Level is displayed.

3. Customize the Performance Service Level by specifying the performance objectives, and then click **Submit** to save the Performance Service Level.

You can apply the new or changed Performance Service Level to workloads (LUNs, NFS File Shares, CIFS Shares) from the Workloads page or when provisioning a new workload.

Managing Storage Efficiency Policies

A Storage Efficiency Policy (SEP) enables you to define the storage efficiency characteristics of a workload. You can assign an SEP to a workload when initially creating the workload, or afterwards by editing the workload.

Storage efficiency includes using technologies, such as thin provisioning, deduplication, and data compression that increase storage utilization and decrease storage costs. While creating SEPs, you can use these space-saving technologies either individually or together to achieve maximum storage efficiency. When you associate the policies with your storage workloads, the specified policy settings are assigned to them. Unified Manager enables you to assign system-defined and user-defined SEPs to optimize storage resources in your data center.

Unified Manager provides two system-defined SEPs: High and Low. These SEPs are applicable to most of the storage workloads in a data center, however, you can create your own policies if the system-defined SEPs do not meet your requirements.

You cannot modify an SEP that is system-defined or that is currently assigned to a workload. You cannot delete an SEP that is assigned to a workload, or if it is the only available SEP.

The Storage Efficiency Policies page lists the available SEPs and enables you to add, edit, and delete customized SEPs. This page displays the following information:

Field	Description
Name	Name of the SEP.
Type	Whether the policy is system-defined or user-defined.
Space Reserve	Whether the volume is thin-provisioned or thick-provisioned.
Deduplication	Whether deduplication is enabled on the workload: <ul style="list-style-type: none">• Inline: Deduplication occurs while being written on the workload• Background: Deduplication occurs in the workload• Disable: Deduplication is disabled on the workload

Field	Description
Compression	<p>Whether data compression is enabled on the workload:</p> <ul style="list-style-type: none"> • Inline: Data compression occurs while being written on the workload • Background: Data compression occurs in the workload • Disable: Data compression is disabled on the workload
Workloads	Number of storage workloads that have been assigned the SEP

Guidelines for creating a custom Storage Efficiency Policy

If the existing SEPs do not meet policy requirements for your storage workloads, you can create a custom SEP. However, it is recommended that you attempt to use the system-defined SEPs for your storage workloads, and only create custom SEPs if necessary.

You can view the SEP assigned to workloads in the All Workloads page and in the Volume / Health details page. You can view the cluster-level data reduction ratio (without Snapshot copies) based on these storage efficiencies in the Capacity panel on the Dashboard and in the Capacity: All Clusters view.

Creating and editing Storage Efficiency Policies

When the system-defined Storage Efficiency Policies do not match your workload requirements, you can create your own Storage Efficiency Policies that are optimized for your workloads.

What you'll need

- You must have the Application Administrator role.
- The Storage Efficiency Policy name must be unique, and you cannot use the following reserved keywords:

High, Low, Unassigned, Learning, Idle, Default, and None.

You create and edit custom Storage Efficiency Policies from the Storage Efficiency Policies page by defining the storage efficiency characteristics you require for the applications that will access storage.



You cannot modify a Storage Efficiency Policy if it is currently assigned to a workload.

Steps

1. In the left navigation pane under **Settings**, select **Policies > Storage Efficiency**.
2. In the **Storage Efficiency Policies** page, click the appropriate button depending on whether you want to create a new Storage Efficiency Policy or if you want to edit an existing Storage Efficiency Policy.

To...	Follow these steps...
Create a new Storage Efficiency Policy	Click Add
Edit an existing Storage Efficiency Policy	Select an existing Storage Efficiency Policy and click Edit

The page to add or edit a Storage Efficiency Policy is displayed.

3. Customize the Storage Efficiency Policy by specifying the storage efficiency characteristics, and then click **Submit** to save the Storage Efficiency Policy.

You can apply the new or changed Storage Efficiency Policy to workloads (LUNs, NFS File Shares, CIFS Shares) from the Workloads page or when provisioning a new workload.

Managing and monitoring MetroCluster configurations

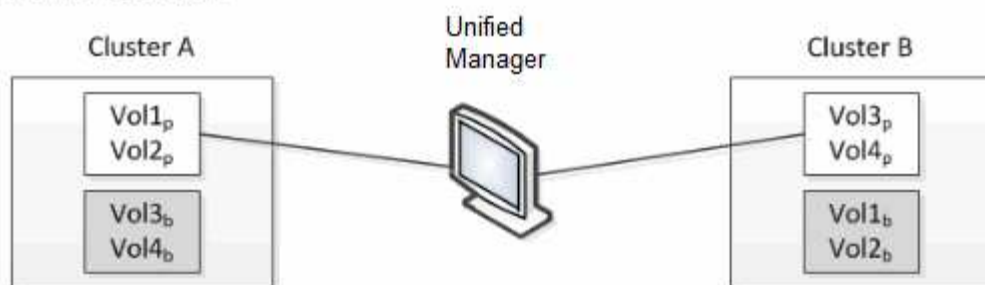
The monitoring support for MetroCluster configurations in the Unified Manager web UI enables you to check for any connectivity issues in your MetroCluster over FC and IP configurations. Discovering a connectivity issue early enables you to manage your MetroCluster configurations effectively.

Volume behavior during switchover and switchback

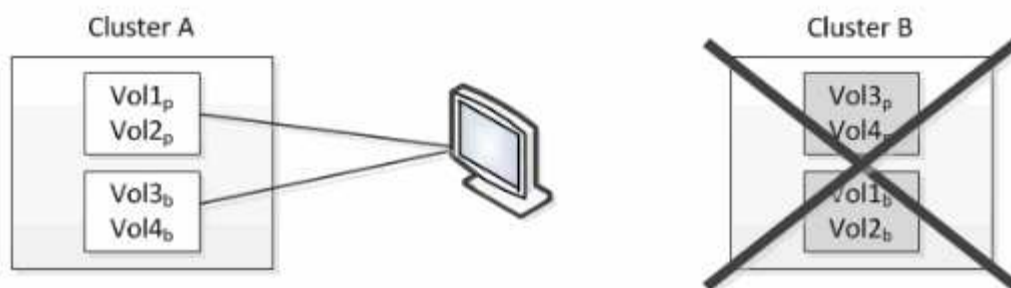
Events that trigger a switchover or switchback cause active volumes to be moved from one cluster to the other cluster in the disaster recovery group. The volumes on the cluster that were active and serving data to clients are stopped, and the volumes on the other cluster are activated and start serving data. Unified Manager monitors only those volumes that are active and running.

Because volumes are moved from one cluster to another, it is recommended that you monitor both clusters. A single instance of Unified Manager can monitor both clusters in a MetroCluster configuration, but sometimes the distance between the two locations necessitates using two Unified Manager instances to monitor both clusters. The following figure shows a single instance of Unified Manager:

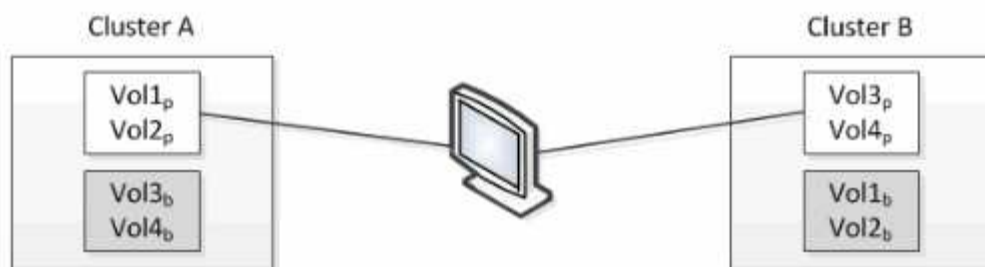
Normal operation




Cluster B fails --- switchover to Cluster A



Cluster B is repaired --- switchback to Cluster B



 = active and monitored

 = inactive and not monitored

The volumes with p in their names indicate the primary volumes, and the volumes with b in their names are mirrored backup volumes that are created by SnapMirror.

During normal operation:

- Cluster A has two active volumes: Vol1p and Vol2p.
- Cluster B has two active volumes: Vol3p and Vol4p.
- Cluster A has two inactive volumes: Vol3b and Vol4b.
- Cluster B has two inactive volumes: Vol1b and Vol2b.

Information pertaining to each of the active volumes (statistics, events, and so on) is collected by Unified Manager. Vol1p and Vol2p statistics are collected by Cluster A, and Vol3p and Vol4p statistics are collected by Cluster B.

After a catastrophic failure causes a switchover of active volumes from Cluster B to Cluster A:

- Cluster A has four active volumes: Vol1p, Vol2p, Vol3b, and Vol4b.

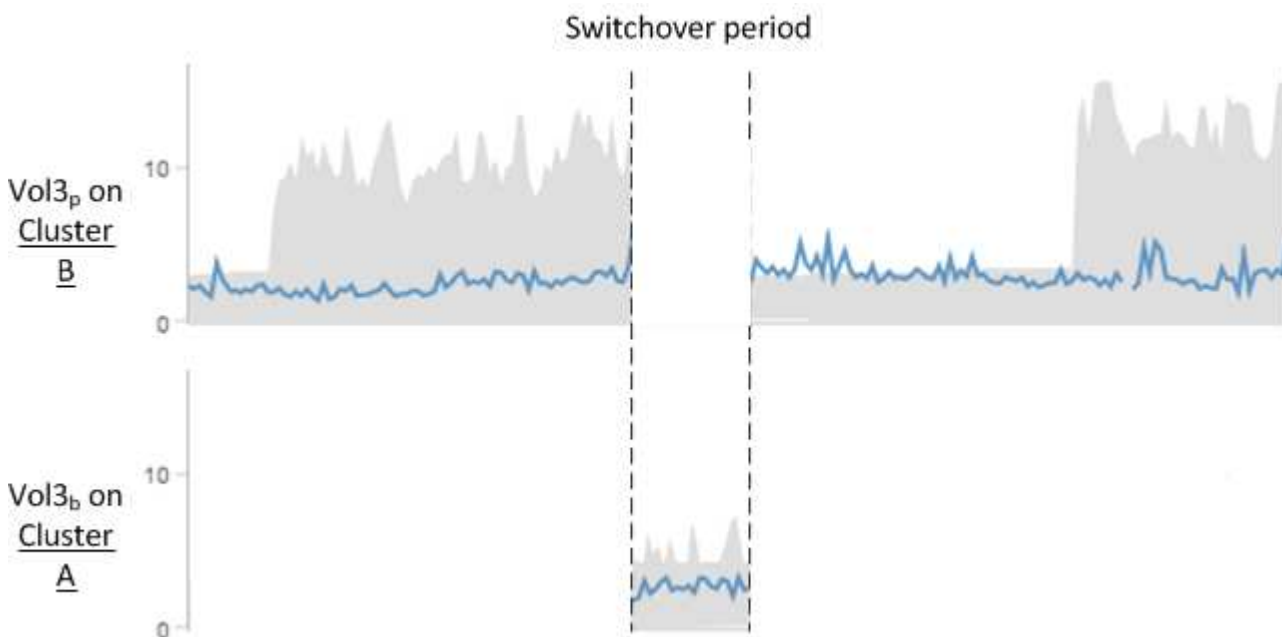
- Cluster B has four inactive volumes: Vol3p, Vol4p, Vol1b, and Vol2b.

As during normal operation, information pertaining to each of the active volumes is collected by Unified Manager. But in this case, Vol1p and Vol2p statistics are collected by Cluster A, and Vol3b and Vol4b statistics are also collected by Cluster A.

Note that Vol3p and Vol3b are not the same volumes, because they are on different clusters. The information in Unified Manager for Vol3p is not the same as Vol3b:

- During switchover to Cluster A, Vol3p statistics and events are not visible.
- On the very first switchover, Vol3b looks like a new volume with no historical information.





When Cluster B is repaired and a switchback is performed, Vol3p is active again on Cluster B, with the historical statistics and a gap of statistics for the period during the switchover. Vol3b is not viewable from Cluster A until another switchover occurs:



- MetroCluster volumes that are inactive, for example, Vol3b on Cluster A after switchback, are identified with the message “This volume was deleted”. The volume is not actually deleted, but it is not currently being monitored by Unified Manager because it is not the active volume.
- If a single Unified Manager is monitoring both clusters in a MetroCluster configuration, volume search returns information for whichever volume is active at that time. For example, a search for "Vol3" would return statistics and events for Vol3b on Cluster A if a switchover has occurred and Vol3 has become active on Cluster A.


Cluster connectivity status definitions for MetroCluster over FC configuration



Connectivity between the clusters in a MetroCluster over FC configuration can be one of the following statuses: Optimal, Impacted, or Down. Understanding the connectivity statuses enables you to manage your MetroCluster configurations effectively.

Connectivity status	Description	Icon displayed
Optimal	Connectivity between the clusters in the MetroCluster configuration is normal.	
Impacted	One or more errors compromise the status of failover availability; however, both of the clusters in the MetroCluster configuration are still up. For example, when the ISL link is down, when the intercluster IP link is down, or when the partner cluster is not reachable.	
Down	Connectivity between the clusters in the MetroCluster configuration is down because one or both of the clusters are down or the clusters are in failover mode. For example, when the partner cluster is down because of a disaster or when there is a planned switchover for testing purposes.	<p>Switchover with errors:</p>  <p>Switchover successful:</p> 

Data mirroring status definitions for MetroCluster over FC

MetroCluster over FC configurations provide data mirroring and the additional ability to initiate a failover if an entire site becomes unavailable. The status of data mirroring between the clusters in a MetroCluster over FC configuration can either be Normal or Mirroring Unavailable. Understanding the status enables you to manage your MetroCluster configurations effectively.

Data mirroring status	Description	Icon displayed
Normal	Data mirroring between the clusters in the MetroCluster configuration is normal.	

Data mirroring status	Description	Icon displayed
Mirroring Unavailable	Data mirroring between the clusters in the MetroCluster configuration is unavailable because of switchover. For example, when the partner cluster is down because of a disaster or when there is a planned switchover for testing purposes.	<p>Switchover with errors:</p>  <p>Switchover successful:</p> 

Monitoring MetroCluster configurations

You can monitor connectivity issues in your MetroCluster configuration. The details include the status of the components and connectivity within a cluster and the connectivity status between the clusters in the MetroCluster configuration. Here, you will get to know how to monitor connectivity issues in clusters protected by MetroCluster over FC and MetroCluster over IP configurations.

You can monitor the MetroCluster configurations from the following views from the Active IQ Unified Manager left navigation pane:

- **Storage > Clusters > Protection: MetroCluster** view
- **Protection > Relationships > Relationship: MetroCluster** view

Unified Manager uses system health alerts to indicate the status of the components and connectivity in the MetroCluster configuration.

What you'll need

- Both the local and remote clusters in a MetroCluster configuration must be added to Active IQ Unified Manager.
- In a MetroCluster over IP configuration, if a Mediator is to be supported, the Mediator should be configured and added to the cluster by corresponding API.
- You must have the Operator, Application Administrator, or Storage Administrator role.

Monitor connectivity issues in MetroCluster over FC configuration

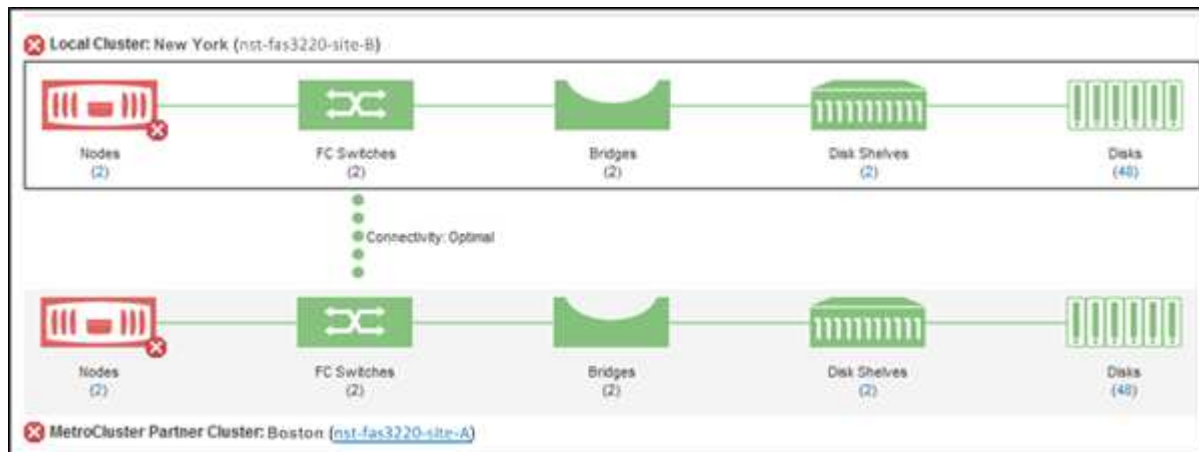
For clusters in a MetroCluster over FC configuration, the connectivity charts are displayed on the **Cluster / Health** details page. Follow these steps.

Steps

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

A list of all of the monitored clusters is displayed.

2. From the **Protection: MetroCluster** view, click the name of the cluster for which you want to view MetroCluster over FC configuration details. Alternately, you can filter by clusters in a MetroCluster configuration.
3. In the **Cluster / Health** details page, click the **MetroCluster Connectivity** tab. The **MetroCluster Connectivity** tab is available for only MetroCluster over FC configurations.



The topology of the MetroCluster configuration is displayed in the corresponding cluster object area. You can use the information displayed in the Cluster / Health details page to rectify any connectivity issues. For example, if the connectivity between the node and the switch in a cluster is down, the following icon is displayed:



If you move the pointer over the icon, you can view detailed information about the generated event.

If you discover connectivity issues in your MetroCluster configuration, you must log in to System Manager or access the ONTAP CLI to resolve the issues.

For more information about determining cluster health, see [Determining cluster health in MetroCluster over FC configuration](#).

Monitor connectivity issues in MetroCluster over IP configuration

For clusters in a MetroCluster over IP configuration, the connectivity charts are displayed on the **Clusters** page. Follow these steps.

Steps

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

A list of all of the monitored clusters is displayed.

2. From the **Protection: MetroClusters** view, click the name of the cluster for which you want to view MetroCluster over IP configuration details. Alternately, you can filter by clusters in a MetroCluster configuration.
3. Expand the row by clicking the caret ∇ icon. The caret icon appears for only a cluster that this protected by MetroCluster over IP configuration.

You can view the topology of the source and mirror sites, as well as the Mediator, if any, used for the connection. You can view the following information:

- Connectivity across the sites
- Health and availability issues, if any, on both the sites
- Mediator-related issues
- Replication related issues.



The following statuses are reported: Critical (❌), Error (⚠️), or Normal (✅). You can also view the aggregate data replication status of the primary and mirror data in the same topology.

In the following diagram, you can see that the intersite connectivity between the source and destination clusters is unavailable, and the Mediator between them is not configured.

4. Click the status icon. A message with the error definition is displayed. If an event has been raised for the issue in your MetroCluster over IP configuration, you can click the **View Event** button on the message and view the event details. When you have resolved the issue and the event, the status icon in this topology turns to normal (✅).
5. You can view further configuration details in the **MetroCluster Overview** and **Protection** sections on the **Configuration** tab of the **Cluster / Health** details page.



Only for a MetroCluster over IP configuration, you can have view the cluster topology on the **Clusters** page. For clusters in a MetroCluster over FC configuration, the topology is displayed on the **MetroCluster Connectivity** tab on the **Cluster / Health** details page.

Related information

- [Cluster / Health details page](#)
- For information about **Relationship:MetroCluster** view, see [Monitoring MetroCluster configurations](#).
- For information about **Relationship: Last 1 month Transfer Status** view, see [Relationship: Last 1 month Transfer Status view](#).
- For information about **Relationship: Last 1 month Transfer Rate** view, see [Relationship: Last 1 month Transfer Rate view](#).
- For information about **Relationship: All Relationships** view, see [Relationship: All Relationships view](#).

Monitoring MetroCluster replication

You can monitor and diagnose the overall health condition of the logical connections while mirroring the data. You can identify the issues or any risk that interrupts mirroring of cluster components such as aggregates, nodes, and storage virtual machines.

Unified Manager uses system health alerts to monitor the status of the components and connectivity in the MetroCluster configuration.

What you'll need

Both the local and remote cluster in the MetroCluster configuration must be added to Unified Manager

Viewing replication for MetroCluster over IP configurations

For MetroCluster over IP configurations, the data replication status is displayed in the topology peek view for a cluster protected by MetroCluster over IP from the following views from the Unified Manager left navigation pane:

- **Storage > Clusters > Protection: MetroCluster** view
- **Protection > Relationships > Relationship: MetroCluster** view

For information, see [Monitor connectivity issues in MetroCluster over IP](#).

Viewing replication for MetroCluster over FC configurations

Follow these steps to determine any issues in the data replication for MetroCluster over FC configuration.

Steps

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

A list of the monitored clusters is displayed.

2. From the **Health: All Clusters** view, click the name of the cluster for which you want to view MetroCluster replication details. On the **Cluster / Health details** page, click the **MetroCluster Replication** tab.

The topology of the MetroCluster configuration to be replicated is displayed at the local site in the corresponding cluster object area with the information about the remote site where the data is being mirrored. If you move the pointer over the icon, you can view detailed information about the generated event.

You can use the information displayed in the Cluster / Health details page to rectify any replication issues. If you discover mirroring issues in your MetroCluster configuration, you must log in to System Manager or access the ONTAP CLI to resolve the issues.

Related information

[Cluster / Health details page](#)

Managing quotas

You can use user and group quotas to limit the amount of disk space or the number of files that a user or a user group can use. You can view user and user group quota information, such as the disk and file usage and the various limits set on disks.

What quota limits are

User quota limits are values that the Unified Manager server uses to evaluate whether space consumption by a user is nearing the limit or has reached the limit that is set by the user's quota. If the soft limit is crossed or if the hard limit is reached, the Unified Manager server generates user quota events.

By default, the Unified Manager server sends a notification email to users who have crossed the quota soft limit or have reached the quota hard limit and for which user quota events are configured. Users with the Application Administrator role can configure alerts that notify the specified recipients of the user or user group

quota events.

You can specify quota limits by using either ONTAP System Manager or the ONTAP CLI.

Viewing user and user group quotas

The Storage VM / Health details page displays information about the user and user group quotas that are configured on the SVM. You can view the name of the user or user group, limits set on the disks and files, used disk and file space, and email address for notification.

What you'll need

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

Steps

1. In the left navigation pane, click **Storage > Storage VMs**.
2. In the **Health: All Storage VMs** view, select a Storage VM and then click the **User and Group Quotas** tab.

Related information

[Adding users](#)

Creating rules to generate email addresses

You can create rules to specify the email address based on the user quota associated with clusters, storage virtual machines (SVMs), volumes, qtrees, users, or user groups. A notification is sent to the specified email address when there is a quota breach.

What you'll need

- You must have the Application Administrator or Storage Administrator role.
- You must have reviewed the guidelines on the Rules to Generate User and Group Quota Email Address page.

You must define the rules for quota email addresses and enter them in the order in which you want to execute them. For example, if you want to use the email address [abc@xyz.com](#) to receive notifications about quota breaches for abc and use the email address dl-\$GROUP@\$DOMAIN for all the other groups, the rules must be listed in the following order:

- if (\$USER == 'abc') then [abc@xyz.com](#)
- if (\$GROUP == *) then dl-\$GROUP@\$DOMAIN

If none of the criteria for the rules you specified are met, then the default rule is used:

```
if ( $USER_OR_GROUP == * ) then $USER_OR_GROUP@$DOMAIN
```

Steps

1. In the left navigation pane, click **General > Quota Email Rules**.
2. Enter the rule based on your criteria.

3. Click **Validate** to validate the syntax of the rule.

An error message is displayed if the syntax of the rule is incorrect. You must correct the syntax and click **Validate** again.

4. Click **Save**.
5. Verify that the email address you created is displayed in the **User and Group Quotas** tab of the Storage VM / Health details page.

Creating an email notification format for user and user group quotas

You can create a notification format for the emails that are sent to a user or a user group when there is a quota-related issue (soft limit breached or hard limit reached).

What you'll need

You must have the Application Administrator or Storage Administrator role.

Steps

1. In the left navigation pane, click **General > Quota Email Format**.
2. Enter or modify the details in the **From**, **Subject**, and **Email Details** fields.
3. Click **Preview** to preview the email notification.
4. Click **Close** to close the preview window.
5. Modify the content of the email notification, if required.
6. Click **Save**.

Editing user and group quota email addresses

You can modify the email addresses based on the user quota associated with clusters, storage virtual machines (SVMs), volumes, qtrees, users, or user groups. You can modify the email address when you want to override the email address generated by rules specified in the Rules to Generate User and Group Quota Email Address dialog box.

What you'll need

- You must have the Operator, Application Administrator, or Storage Administrator role.
- You must have reviewed the [guidelines for creating rules](#).

If you edit an email address, the rules to generate the user and group quota email addresses are no longer applicable to the quota. For notifications to be sent to the email address generated by the rules specified, you must delete the email address and save the change.

Steps

1. In the left navigation pane, click **Storage > SVMs**.
2. In the **Health: All Storage VMs** view, select an SVM and then click the **User and Group Quotas** tab.
3. Click **Edit Email Address** below the row of tabs.
4. In the **Edit Email Address** dialog box, perform the appropriate action:

If...	Then...
You want notifications to be sent to the email address generated by the rules specified	<ol style="list-style-type: none"> Delete the email address in the Email Address field. Click Save. Refresh the browser by pressing F5 to reload the Edit Email Address dialog box. The email address generated by the specified rule is displayed in the Email Address field.
You want notifications to be sent to a specified email address	<ol style="list-style-type: none"> Modify the email address in the Email Address field. Click Save. The rules to generate the user and group quota email addresses are no longer applicable to the quota.

Understanding more about quotas

Understanding the concepts about quotas helps you to manage your user quotas and user group quotas efficiently.

Overview of the quota process

Quotas can be soft or hard. Soft quotas cause ONTAP to send a notification when specified limits are exceeded, and hard quotas prevent a write operation from succeeding when specified limits are exceeded.

When ONTAP receives a request from a user or user group to write to a FlexVol volume, it checks to see whether quotas are activated on that volume for the user or user group and determines the following:

- Whether the hard limit will be reached

If yes, the write operation fails when the hard limit is reached and the hard quota notification is sent.

- Whether the soft limit will be breached

If yes, the write operation succeeds when the soft limit is breached and the soft quota notification is sent.

- Whether a write operation will not exceed the soft limit

If yes, the write operation succeeds and no notification is sent.

About quotas

Quotas provide a way to restrict or track the disk space and number of files used by a user, group, or qtree. You specify quotas using the `/etc/quotas` file. Quotas are applied to a specific volume or qtree.

Why you use quotas

You can use quotas to limit resource usage in FlexVol volumes, to provide notification when resource usage reaches specific levels, or to track resource usage.

You specify a quota for the following reasons:

- To limit the amount of disk space or the number of files that can be used by a user or group, or that can be contained by a qtree
- To track the amount of disk space or the number of files used by a user, group, or qtree, without imposing a limit
- To warn users when their disk usage or file usage is high

Description of quotas dialog boxes

You can use the appropriate option in the User and Group Quotas tab in the Health: All Storage VMs view to configure the format of the email notification that is sent when a quota-related issue occurs and to configure rules to specify email addresses based on the user quota.

Email Notification Format page

The Email Notification Format page displays the rules of the email that is sent to a user or a user group when there is a quota-related issue (soft limit breached or hard limit reached).

The email notification is sent only when the following user or user group quota events are generated: User or Group Quota Disk Space Soft Limit Breached, User or Group Quota File Count Soft Limit Breached, User or Group Quota Disk Space Hard Limit Reached, or User or Group Quota File Count Hard Limit Reached.

- **From**

Displays the email address from which the email is sent, which you can modify. By default, this is the email address that is specified Notifications page.

- **Subject**

Displays the subject of the notification email.

- **Email Details**

Displays the text of the notification email. You can modify the text based on your requirements. For example, you can provide information related to the quota attributes and reduce the number of keywords. However, you should not modify the keywords.

Valid keywords are as follows:

- `$EVENT_NAME`

Specifies the event name that caused the email notification.

- `$QUOTA_TARGET`

Specifies the qtree or volume on which the quota is applicable.

- **\$QUOTA_USED_PERCENT**

Specifies the percentage of disk hard limit, disk soft limit, file hard limit, or file soft limit that is used by the user or user group.

- **\$QUOTA_LIMIT**

Specifies the disk hard limit or file hard limit that is reached by the user or user group and one of the following events is generated:

- User or Group Quota Disk Space Hard Limit Reached
- User or Group Quota Disk Space Soft Limit Reached
- User or Group Quota File Count Hard Limit Reached
- User or Group Quota File Count Soft Limit Reached

- **\$QUOTA_USED**

Specifies the disk space used or the number of files created by the user or user group.

- **\$QUOTA_USER**

Specifies the user or user group name.

Command buttons

The command buttons enable you to preview, save, or cancel the changes made to the email notification format:

- **Preview**

Displays a preview of the notification email.

- **Restore to Factory Defaults**

Enables you to restore the notification format to the factory default values.

- **Save**

Saves the changes made to the notification format.

Rules to Generate User and Group Quota Email Address page

The Rules to Generate User and Group Quota Email Address page enables you to create rules to specify email addresses based on the user quota associated with clusters, SVMs, volumes, qtrees, users, or user groups. A notification is sent to the specified email address when a quota is breached.

Rules area

You must define the rules for a quota email address. You can also add comments to explain the rules.

How you define rules

You must enter the rules in the order in which you want to execute them. If the first rule's criterion is met, then the email address is generated based on this rule. If the criterion is not met, then the criterion for the next rule is considered, and so on. Each line lists a separate rule. The default rule is the last rule in the list. You can change the priority order of rules. However, you cannot change the order of the default rule.

For example, if you want to use the email address qtree1@xyz.com to receive notifications about quota breaches for qtree1 and use the email address admin@xyz.com for all the other qtrees, the rules must be listed in the following order:

- if (\$QTREE == 'qtree1') then qtree1@xyz.com
- if (\$QTREE == *) then admin@xyz.com

If none of the criteria for the rules you specified are met, then the default rule is used:

```
if ( $USER_OR_GROUP == * ) then $USER_OR_GROUP@$DOMAIN
```

If more than one user has the same quota, the names of the users are displayed as comma-separated values and the rules are not applicable for the quota.

How you add comments

You can add comments to explain the rules. You should use # at the start of each comment and each line lists a separate comment.

Rules syntax

The syntax of the rule must be one of the following:

- if (valid variableoperator *) then email ID@domain name

if is a keyword and is in lowercase. The operator is ==. The email ID can contain any character, the valid variables \$USER_OR_GROUP, \$USER, or \$GROUP, or a combination of any character and the valid variables \$USER_OR_GROUP, \$USER, or \$GROUP. The domain name can contain any character, the valid variable \$DOMAIN, or a combination of any character and the valid variable \$DOMAIN. Valid variables can be in uppercase or lowercase but must not be a combination of both. For example, \$domain and \$DOMAIN are valid, but \$Domain is not a valid variable.

- if (valid variableoperator `string`) then email ID@domain name

if is a keyword and is lowercase. The operator can be contains or ==. The email ID can contain any character, the valid variables \$USER_OR_GROUP, \$USER, or \$GROUP, or a combination of any character and the valid variables \$USER_OR_GROUP, \$USER, or \$GROUP. The domain name can contain any character, the valid variable \$DOMAIN, or a combination of any character and the valid variable \$DOMAIN. Valid variables can be in uppercase or lowercase but must not be a combination of both. For example, \$domain and \$DOMAIN are valid, but \$Domain is not a valid variable.

Command buttons

The command buttons enable you to save, validate, or cancel the created rules:

- **Validate**

Validates the syntax of the created rule. If there are errors during validation, the rule that generates the

error is displayed along with an error message.

- **Restore to Factory Defaults**

Enables you to restore the address rules to the factory default values.

- **Save**

Validates the syntax of the rule and saves the rule if there are no errors. If there are errors during validation, the rule that generates the error is displayed along with an error message.

Troubleshooting

Troubleshooting information helps you to identify and resolve issues you encounter when using Unified Manager.

Adding disk space to the Unified Manager database directory

The Unified Manager database directory contains all of the health and performance data collected from ONTAP systems. Some circumstances may require that you increase the size of the database directory.

For example, the database directory may get full if Unified Manager is collecting data from a large number of clusters where each cluster has many nodes. You will receive a warning event when the database directory is 90% full, and a critical event when the directory is 95% full.



No additional data is collected from clusters after the directory reaches 95% full.

The steps required to add capacity to the data directory are different depending on whether Unified Manager is running on a VMware ESXi server, on a Red Hat or CentOS Linux server, or on a Microsoft Windows server.

Adding space to the data disk of the VMware virtual machine

If you need to increase the amount of space on the data disk for the Unified Manager database, you can add capacity after installation by increasing disk space using the Unified Manager maintenance console.

What you'll need

- You must have access to the vSphere Client.
- The virtual machine must have no snapshots stored locally.
- You must have the maintenance user credentials.

We recommend that you back up your virtual machine before increasing the size of virtual disks.

Steps

1. In the vSphere client, select the Unified Manager virtual machine, and then add more disk capacity to data disk 3. See the VMware documentation for details.

In some rare cases the Unified Manager deployment uses "Hard Disk 2" for the data disk instead of "Hard

Disk 3". If this has occurred in your deployment, increase the space of whichever disk is larger. The data disk will always have more space than the other disk.

2. In the vSphere client, select the Unified Manager virtual machine, and then select the **Console** tab.
3. Click in the console window, and then log in to the maintenance console using your user name and password.
4. In the **Main Menu**, enter the number for the **System Configuration** option.
5. In the **System Configuration Menu**, enter the number for the **Increase Data Disk Size** option.

Adding space to the data directory of the Linux host

If you allotted insufficient disk space to the `/opt/netapp/data` directory to support Unified Manager when you originally set up the Linux host and then installed Unified Manager, you can add disk space after installation by increasing disk space on the `/opt/netapp/data` directory.

What you'll need

You must have root user access to the Red Hat Enterprise Linux or CentOS Linux machine on which Unified Manager is installed.

We recommend that you back up the Unified Manager database before increasing the size of the data directory.

Steps

1. Log in as root user to the Linux machine on which you want to add disk space.
2. Stop the Unified Manager service and the associated MySQL software in the order shown: `systemctl stop ocieau ocie mysqld`
3. Create a temporary backup folder (for example, `/backup-data`) with sufficient disk space to contain the data in the current `/opt/netapp/data` directory.
4. Copy the content and privilege configuration of the existing `/opt/netapp/data` directory to the backup data directory:

```
cp -arp /opt/netapp/data/* /backup-data
```

5. If SE Linux is enabled:
 - a. Get the SE Linux type for folders on existing `/opt/netapp/data` folder:

```
se_type=`ls -Z /opt/netapp/data | awk '{print $4}' | awk -F: '{print $3}' | head -1
```

The system returns a confirmation similar to the following:

```
echo $se_type
mysqld_db_t
```

- b. Run the `chcon` command to set the SE Linux type for the backup directory:

```
chcon -R --type=mysqlld_db_t /backup-data
```

6. Remove the contents of the `/opt/netapp/data` directory:

a. `cd /opt/netapp/data`

b. `rm -rf *`

7. Expand the size of the `/opt/netapp/data` directory to a minimum of 150 GB through LVM commands or by adding extra disks.



If you have created `/opt/netapp/data` from a disk, then you should not try to mount `/opt/netapp/data` as an NFS or CIFS share. Because, in this case, if you try to expand the disk space, some LVM commands, such as `resize` and `extend` might not work as expected.

8. Confirm that the `/opt/netapp/data` directory owner (mysql) and group (root) are unchanged:

```
ls -ltr /opt/netapp/ | grep data
```

The system returns a confirmation similar to the following:

```
drwxr-xr-x. 17 mysql root 4096 Aug 28 13:08 data
```

9. If SE Linux is enabled, confirm that the context for the `/opt/netapp/data` directory is still set to `mysqlld_db_t`:

a. `touch /opt/netapp/data/abc`

b. `ls -Z /opt/netapp/data/abc`

The system returns a confirmation similar to the following:

```
-rw-r--r--. root root unconfined_u:object_r:mysqlld_db_t:s0  
/opt/netapp/data/abc
```

10. Delete the file `abc` so that this extraneous file does not cause a database error in the future.

11. Copy the contents from `backup-data` back to the expanded `/opt/netapp/data` directory:

```
cp -arp /backup-data/* /opt/netapp/data/
```

12. If SE Linux is enabled, run the following command:

```
chcon -R --type=mysqlld_db_t /opt/netapp/data
```

13. Start the MySQL service:

```
systemctl start mysqld
```

14. After the MySQL service is started, start the `ocie` and `ocieau` services in the order shown:

```
systemctl start ocie ocieau
```

15. After all of the services are started, delete the backup folder `/backup-data`:

```
rm -rf /backup-data
```

Adding space to the logical drive of the Microsoft Windows server

If you need to increase the amount of disk space for the Unified Manager database, you can add capacity to the logical drive on which Unified Manager is installed.

What you'll need

You must have Windows administrator privileges.

We recommend that you back up the Unified Manager database before adding disk space.

Steps

1. Log in as administrator to the Windows server on which you want to add disk space.
2. Follow the step that corresponds to method you want to use to add more space:

Option	Description
On a physical server, add capacity to the logical drive on which the Unified Manager server is installed.	Follow the steps in the Microsoft topic: Extend a Basic Volume
On a physical server, add a hard disk drive.	Follow the steps in the Microsoft topic: Adding Hard Disk Drives
On a virtual machine, increase the size of a disk partition.	Follow the steps in the VMware topic: Increasing the size of a disk partition

Changing the performance statistics collection interval

The default collection interval for performance statistics is 5 minutes. You can change this interval to 10 or 15 minutes if you find that collections from large clusters are not finishing within the default time. This setting affects the collection of statistics from all clusters that this instance of Unified Manager is monitoring.

What you'll need

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

The issue of performance statistics collections not finishing on time is indicated by the banner messages `Unable to consistently collect from cluster <cluster_name>` or `Data collection is taking too long on cluster <cluster_name>`.

You should change the collection interval only when required because of a statistics collections issue. Do not change this setting for any other reason.



Changing this value from the default setting of 5 minutes can affect the number and frequency of performance events that Unified Manager reports. For example, system-defined performance thresholds trigger events when the policy is exceeded for 30 minutes. When using 5-minute collections, the policy must be exceeded for six consecutive collections. For 15-minute collections the policy must be exceeded for only two collection periods.

A message at the bottom of the Cluster Setup page indicates the current statistical data collection interval.

Steps

1. Log in using SSH as the maintenance user to the Unified Manager host.

The Unified Manager maintenance console prompts are displayed.

2. Type the number of the menu option labeled **Performance Polling Interval Configuration**, and then press Enter.
3. If prompted, enter the maintenance user password again.
4. Type the number for the new polling interval that you want to set, and then press Enter.

If you changed the Unified Manager collection interval to 10 or 15 minutes, and you have a current connection to an external data provider (such as Graphite), you must change the data provider transmit interval so that it is equal to, or greater, than the Unified Manager collection interval.

Changing the length of time Unified Manager retains event and performance data

By default, Unified Manager stores event data and performance data for 6 months for all monitored clusters. After this time, older data is automatically deleted to make room for new data. This default timeframe works well for most configurations, but very large configurations with many clusters and nodes may need to reduce the retention period so that Unified Manager operates optimally.

What you'll need

You must have the Application Administrator role.

You can change the retention periods for these two types of data in the Data Retention page. These settings affect the retention of data from all clusters that this instance of Unified Manager is monitoring.



Unified Manager collects performance statistics every 5 minutes. Each day the 5-minute statistics are summarized into hourly performance statistics. It retains 30 days of 5-minute historical performance data and 6 months of hourly summarized performance data (by default).

You should reduce the retention period only if you are running out of space or if backup and other operations are taking a very long time to complete. Reducing the retention period has the following effects:

- Old performance data is deleted from the Unified Manager database after midnight.
- Old event data is deleted from the Unified Manager database immediately.
- Events prior to the retention period will no longer be available to view in the user interface.

- Locations in the UI where hourly performance statistics are displayed will be blank prior to the retention period.
- If the event retention period exceeds the performance data retention period, a message will be displayed under the performance slider warning that older performance events may not have backing data in their associated charts.

Steps

1. In the left navigation pane, click **Policies > Data Retention**.
2. In the **Data Retention** page, select the slider tool in the Event Retention or Performance Data Retention area and move it to the number of months that data should be retained, and click **Save**.

Unknown authentication error

When you are performing an authentication-related operation such as adding, editing, deleting, or testing remote users or groups, the following error message might be displayed: Unknown authentication error.

Cause

This problem can occur if you have set an incorrect value for the following options:

- Administrator Name of the Active Directory authentication service
- Bind Distinguished Name of the OpenLDAP authentication service

Corrective action

- a. In the left navigation pane, click **General > Remote Authentication**.
- b. Based on the authentication service that you have selected, enter the appropriate information for Administrator Name or Bind Distinguished Name.
- c. Click **Test Authentication** to test the authentication with the details that you specified.
- d. Click **Save**.

User not found

When you are performing an authentication-related operation such as adding, editing, deleting, or testing remote users or groups, the following error message is displayed: User not found.

Cause

This problem can occur if the user exists in the AD server or LDAP server, and if you have set the base distinguished name to an incorrect value.

Corrective action

- a. In the left navigation pane, click **General > Remote Authentication**.
- b. Enter the appropriate information for base distinguished name.
- c. Click **Save**.

Issue with adding LDAP using Other authentication services

When you select Others as the Authentication service, the user and group Object Class retain the values from the previously selected template. If the LDAP server does not use the same values, the operation might fail.

Cause

The users are not configured correctly in OpenLDAP.

Corrective action

You can manually fix this issue by using one of the following workarounds.

If your LDAP user object class and group object class are user and group, respectively, perform the following steps:

- a. In the left navigation pane, click **General > Remote Authentication**.
- b. In the **Authentication Service** drop-down menu, select **Active Directory**, and then select **Others**.
- c. Complete the text fields.

If your LDAP user object class and group object class are posixAccount and posixGroup, respectively, perform the following steps:

- a. In the left navigation pane, click **General > Remote Authentication**.
- b. In the **Authentication Service** drop-down menu, select **OpenLDAP**, and then select **Others**.
- c. Complete the text fields.

If the first two workarounds do not apply, call the `option-set` API, and set the `auth.ldap.userObjectClass` and `auth.ldap.groupObjectClass` options to the correct values.

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.